

ANALYZING TIMED SYSTEMS USING TREE AUTOMATA

S. AKSHAY, PAUL GASTIN, AND SHANKARA NARAYANAN KRISHNA

Dept. of CSE, IIT Bombay, Powai, Mumbai 400076, India
e-mail address: akshayss@cse.iitb.ac.in

LSV, ENS-Cachan, CNRS, Université Paris-Saclay, 94235 Cachan, France
e-mail address: paul.gastin@lsv.ens-cachan.fr

Dept. of CSE, IIT Bombay, Powai, Mumbai 400076, India
e-mail address: krishnas@cse.iitb.ac.in

ABSTRACT. Timed systems, such as timed automata, are usually analyzed using their operational semantics on timed words. The classical region abstraction for timed automata reduces them to (untimed) finite state automata with the same time-abstract properties, such as state reachability. We propose a new technique to analyze such timed systems using finite tree automata instead of finite word automata. The main idea is to consider timed behaviors as graphs with matching edges capturing timing constraints. Such graphs can be interpreted in trees opening the way to tree automata based techniques which are more powerful than analysis based on word automata. The technique is quite general and applies to many timed systems. In this paper, as an example, we develop the technique on timed pushdown systems, which have recently received considerable attention. Further, we also demonstrate how we can use it on timed automata and timed multi-stack pushdown systems (with boundedness restrictions).

1. INTRODUCTION

The advent of timed automata [4] marked the beginning of an era in the verification of real-time systems. Today, timed automata form one of the well accepted real-time modelling formalisms, using real-valued variables called clocks to capture time constraints. The decidability of the emptiness problem for timed automata is achieved using the notion of region abstraction. This gives a sound and finite abstraction of an infinite state system, and has paved the way for state-of-the-art tools like UPPAAL [6], which have successfully been used in the verification of several complex timed systems. In recent times [1, 7, 13] there has been a lot of interest in the theory of verification of more complex timed systems enriched with features such as concurrency, communication between components and recursion with single or multiple threads. In most of these approaches, decidability has been obtained by cleverly extending the fundamental idea of region or zone abstractions.

1998 ACM Subject Classification: F.1.1 Models of Computation.

Key words and phrases: Timed automata, tree automata, pushdown systems, tree-width.

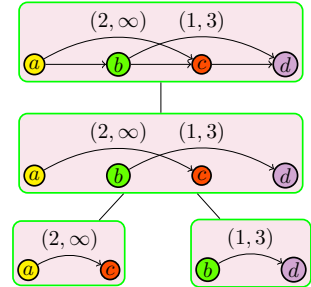
In this paper, we give a technique for analyzing timed systems, inspired from a different approach based on graphs and tree automata. This approach has been exploited for analyzing various types of *untimed systems*, e.g., [17, 10]. The basic template of this approach has three steps: (1) capture the behaviors of the system as graphs, (2) show that the class of graphs that are actual behaviors of the system is MSO-definable, and (3) show that this class of graphs has bounded tree-width (or clique-width or split-width), or restrict the analysis to such bounded behaviors. Then, non-emptiness of the given system boils down to the satisfiability of an MSO sentence on graphs of bounded tree-width, which is decidable by Courcelle’s theorem. But, by providing a direct construction of the tree automaton, it is possible to obtain a good complexity for the decision procedure.

We lift this technique to deal with timed systems by abstracting timed word behaviors of timed systems as graphs consisting of untimed words with additional time-constraint edges, called words with timing constraints (TCWs). The main complication here is that a TCW describes an abstract run of the timed system, where the constraints are recorded but not checked. The TCW corresponds to an actual concrete run if and only if it is *realizable*, i.e., we can find time-stamps realizing the TCW. Thus, we are interested in the class of graphs which are *realizable* TCWs.

For this class of graphs, the above template tells us that we need to show (i) these graphs have a bounded tree-width and (ii) the property of being a realizable TCW is MSO-defineable. Then by Courcelle’s theorem we obtain a tree automaton accepting this. However, as mentioned earlier, the MSO to tree-automaton approach does not give a good complexity in terms of size of the tree automaton. To obtain an optimal complexity, instead of going via Courcelle’s theorem, we directly build a tree automaton. Using tree decompositions of graph behaviours having bounded split/tree-width and constructing tree automata proved to be a very successful technique for the analysis of *untimed* infinite state systems [17, 11, 10, 2]. This paper opens up this powerful technique for analysis of *timed* systems.

Thus our contributions are the following. We start by showing that behaviors of timed systems can be written as words with *simple* timing constraints (STCWs), i.e., words where each position has at most one timing constraint (incoming or outgoing) attached to it. This is done by breaking each transition of the timed system into a sequence of “micro-transitions”, so that at each micro-transition, only one timing constraint is attached.

Next, we show that STCWs that arise as behaviours of certain classes of timed systems (e.g., timed automata or timed pushdown systems) are graphs of bounded *special* tree-width. Special tree-width is a graph complexity measure, arising out of a special tree decomposition of a graph, as introduced by Courcelle in [8]. To establish the bound, we play a so-called *split-game*, which gives a bound on what is called the *split-width*, a notion that was introduced for graph behaviours of untimed systems and which has proven to be very useful for untimed systems [11, 10, 2]. Establishing a relationship between the split-width and special tree-width, we obtain a bound on special tree-width (which also implies a bound on general tree-width as shown in [8]). As a result of this bound, we infer that our graphs admit binary tree decompositions as depicted in the adjoining figure. Each node of the tree decomposition depicts a partial behaviour of the system in a bounded manner. By combining these behaviours as we go up the tree, we obtain a full behaviour of the system.



Our final and most technically challenging step is to construct tree automata that works on such tree decompositions and accepts only those whose roots are labeled by *realizable* STCWs generated by the timed system \mathcal{S} . Thus, checking non-emptiness of the timed system reduces to checking non-emptiness for this tree automaton (which is PTIME in size of the tree automaton). The construction of this tree automaton is done in two phases. First, given the bound on special tree-width of the graph behaviours of \mathcal{S} , we construct a tree automaton that accepts all such trees whose roots are labeled by realizable STCWs with respect to the maximal constant given by the system, and whose nodes are all bounded. For this, we need to check that the graphs at the root are indeed words with valid timing constraints and there exists a time-stamping that realizes the STCW. If we could maintain the partially constructed STCW in a state of the tree automaton along with a guess of the time-stamp at each vertex, we could easily check this. However, the tree automaton has only finitely many states, so while processing the tree bottom-up, we need a finite abstraction of the STCW which remembers only finitely many time values. We show that this is indeed possible by coming up with an abstraction where it is sufficient to remember the modulo M values of time stamps, where M is one more than the maximal constant that appears in the timed system.

In the second phase, we refine this tree automaton to obtain another tree automaton that only accepts those trees that are generated by the system. Yet again, the difficulty is to ensure correct matching of (i) clock constraints between points where a clock is reset and a constraint is checked, (ii) push and pop transitions by keeping only finite amount of information in the states of the tree automaton. Once both of these are done, the final tree automaton satisfies all our constraints.

To illustrate the technique, we have reproved the decidability of non-emptiness of timed automata and timed pushdown automata (TPDA), by showing that both these models have a split-width ($|X| + 4$ and $4|X| + 6$) that is linear in the number of clocks, $|X|$, of the given timed system. This bound directly tells us the amount of information that we need to maintain in the construction of the tree automata. For TPDA we obtain an EXPTIME algorithm, matching the known lower-bound for the emptiness problem of TPDA. For timed automata, since the split-trees are word-like (at each binary node, one subtree is small) we may use word automata instead of tree automata, reducing the complexity from EXPTIME to PSPACE, again matching the lower-bound. Interestingly, if one considers TPDA with no explicit clocks, but the stack is timed, then the split-width is a constant, 2. In this case, we have a polynomial time procedure to decide emptiness, assuming a unary encoding of constants in the system. To further demonstrate the power of our technique, we derive a new decidability result for non-emptiness of timed multi-stack pushdown automata under bounded rounds, by showing that the split-width of this model is again linear in the number of clocks, stacks and rounds. Exploring decidable subclasses of untimed multi-stack pushdown systems is a very active research area [5, 12, 14, 16, 15], and our technique can extend these to handle time.

It should be noticed that the tree automaton for validity/realizability (the most involved construction) is independent of the timed system under study. Hence, to apply the technique to other systems, one only needs to prove the bound on split-width and to show that their runs can be captured by tree automata. This is a major difference compared to many existing techniques for timed systems which are highly system dependent. Finally, we mention an orthogonal approach to deal with timed systems given in [7], where the authors

show the decidability of the non-emptiness problem for a class of timed pushdown automata by reasoning about sets with timed-atoms.

An extended abstract of this paper was presented in [3]. There are however, significant differences from that version as we detail now. First, the technique used to prove the main theorem of building a tree automaton to check realizability is completely different. In [3], we showed an automaton which checks for non-existence of negative weight cycles (which implies the existence of a realizable time-stamping). This required a rather complicated proof to show that the constants can be bounded. In fact, we first build an infinite state tree automaton and then show that we can get a finite state abstraction for it. In contrast, our proof in this article directly builds a tree automaton that checks for existence of time-stamps realizing a run. This allows us to improve the complexity and also gives a less involved proof. Second, we complete the proof details and compute the complexity for tree automata for multistack pushdown systems with bounded round restriction, which was announced in [3]. All sketches from the earlier version have been replaced and enhanced by rigorous proofs in this article. Further, several supporting examples and intuitive explanations have been added throughout to aid in understanding.

2. GRAPHS FOR BEHAVIORS OF TIMED SYSTEMS

We fix an alphabet Σ and use Σ_ε to denote $\Sigma \cup \{\varepsilon\}$ where ε is the silent action. For an integer M , we also fix $\mathcal{I}(M)$, to be a finite set of closed intervals whose endpoints are integers between 0 and M , and which contains the special interval $[0, 0]$. When M is irrelevant or clear from the context, we just write \mathcal{I} instead of $\mathcal{I}(M)$. For a set S , we use $\leq \subseteq S \times S$ to denote a partial or total order on S . For any $x, y \in S$, we write $x < y$ if $x \leq y$ and $x \neq y$, and $x \prec y$ if $x < y$ and there does not exist $z \in S$ such that $x < z < y$.

2.1. Preliminaries: Timed words and timed (pushdown) automata. An ε -timed word is a sequence $w = (a_1, t_1) \dots (a_n, t_n)$ with $a_1 \dots a_n \in \Sigma_\varepsilon^+$ and $(t_i)_{1 \leq i \leq n}$ is a non-decreasing sequence of real time values. If $a_i \neq \varepsilon$ for all $1 \leq i \leq n$, then w is a *timed word*. The projection on Σ of an ε -timed word is the timed word obtained by removing ε -labelled positions. We define the two basic system models that we consider in this article.

Dense-timed pushdown automata (TPDA), introduced in [1], are an extension of timed automata [4], and operate on a finite set of real-valued clocks and a stack which holds symbols with their ages. The age of a symbol in the stack represents time elapsed since it was pushed on to the stack. Formally, a TPDA \mathcal{S} is a tuple $(S, s_0, \Sigma, \Gamma, \Delta, X, F)$ where S is a finite set of states, $s_0 \in S$ is the initial state, Σ, Γ , are respectively a finite set of input, stack symbols, Δ is a finite set of transitions, X is a finite set of real-valued variables called clocks, $F \subseteq S$ are final states. A transition $t \in \Delta$ is a tuple $(s, \gamma, a, \text{op}, R, s')$ where $s, s' \in S$, $a \in \Sigma$, γ is a finite conjunction of atomic formulae of the kind $x \in I$ for $x \in X$ and $I \in \mathcal{I}$, $R \subseteq X$ are the clocks reset, op is one of the following stack operations:

- (1) **nop** does not change the contents of the stack,
- (2) \downarrow_c where $c \in \Gamma$ is a push operation that adds c on top of the stack, with age 0.
- (3) \uparrow_c^I where $c \in \Gamma$ is a stack symbol and $I \in \mathcal{I}$ is an interval, is a pop operation that removes the top most symbol of the stack if it is a c with age in the interval I .

Timed automata (TA) can be seen as TPDA using **nop** operations only. This definition of TPDA is equivalent to the one in [1], but allows checking conjunctive constraints and stack operations together. In [7], it is shown that TPDA of [1] are expressively equivalent to

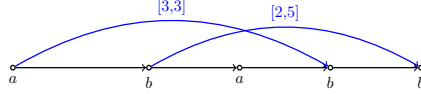
timed automata with an untimed stack. Nevertheless, our technique is oblivious to whether the stack is timed or not, hence we focus on the syntactically more succinct model of TPDA with timed stacks and get good complexity bounds.

The operational semantics of TPDA and TA can be given in terms of timed words and we refer to [1] for the formal definition. Instead, we are interested in an alternate yet equivalent semantics for TPDA using graphs with timing constraints, that we define next.

2.2. Abstractions of timed behaviors as graphs.

Definition 2.1. A word with timing constraints (TCW) over Σ, \mathcal{I} is a structure $\mathcal{V} = (V, \rightarrow, (\curvearrowright^I)_{I \in \mathcal{I}}, \lambda)$ where V is a finite set of positions (also referred to as vertices or points), $\lambda: V \rightarrow \Sigma_\epsilon$ labels each position, the reflexive transitive closure $\leq = \rightarrow^*$ is a total order on V and $\rightarrow = <$ is the successor relation, $\curvearrowright^I \subseteq < = \rightarrow^+$ gives the pairs of positions carrying a timing constraint associated with the interval I .

For any position $i \in V$, the *indegree* (resp. *outdegree*) of i is the number of positions j such that $(j, i) \in \curvearrowright$ (resp. $(i, j) \in \curvearrowright$). A TCW is *simple* (denoted STCW) if each position has at most one timing constraint (incoming or outgoing) attached to it, i.e., for all $i \in V$, $\text{indegree}(i) + \text{outdegree}(i) \leq 1$. A TCW is depicted below with positions $1, 2, \dots, 5$ labelled over $\{a, b\}$. $\text{indegree}(4)=1$, $\text{outdegree}(1)=1$ and $\text{indegree}(3)=0$. The curved edges decorated with intervals connect the positions related by \curvearrowright , while straight edges are the successor relation \rightarrow . Note that this TCW is *simple*.



Consider a TCW $\mathcal{V} = (V, \rightarrow, (\curvearrowright^I)_{I \in \mathcal{I}}, \lambda)$ with $V = \{1, \dots, n\}$. A timed word w is a *realization* of \mathcal{V} if it is the projection on Σ of an ϵ -timed word $w' = (\lambda(1), t_1) \dots (\lambda(n), t_n)$ such that $t_j - t_i \in I$ for all $(i, j) \in \curvearrowright^I$. In other words, a TCW is realizable if there exists a timed word w which is a realization of \mathcal{V} . For example, the timed word $(a, 0.9)(b, 2.1)(a, 2.1)(b, 3.9)(b, 5)$ is a realization of the TCW depicted above, while $(a, 1.2)(b, 2.1)(a, 2.1)(b, 3.9)(b, 5)$ is not.

2.3. Semantics of TPDA (and TA) as simple TCWs. We define the semantics in terms of simple TCWs. An STCW $\mathcal{V} = (V, \rightarrow, (\curvearrowright^I)_{I \in \mathcal{I}}, \lambda)$ is said to be generated or accepted by a TPDA \mathcal{S} if there is an accepting abstract run $\rho = (s_0, \gamma_1, a_1, \text{op}_1, R_1, s_1)(s_1, \gamma_2, a_2, \text{op}_2, R_2, s_2) \dots (s_{n-1}, \gamma_n, a_n, \text{op}_n, R_n, s_n)$ of \mathcal{S} such that, $s_n \in F$ and

- the sequence of push-pop operations is well-nested: in each prefix $\text{op}_1 \dots \text{op}_k$ with $1 \leq k \leq n$, number of pops is at most number of pushes, and in the full sequence $\text{op}_1 \dots \text{op}_n$, they are equal.
- We have $V = V_0 \uplus V_1 \uplus \dots \uplus V_n$ with $V_i \times V_j \subseteq \rightarrow^+$ for $0 \leq i < j \leq n$. Each transition $\delta_i = (s_{i-1}, \gamma_i, a_i, \text{op}_i, R_i, s_i)$ gives rise to a sequence of consecutive points V_i in the STCW. The transition δ_i is simulated by a sequence of “micro-transitions” as depicted below (left) and it represents an STCW shown below (right). Incoming red edges check guards from γ_i (wrt different clocks) while outgoing green edges depict resets from R_i that will be checked later. Further, the outgoing edge on the central node labeled a_i represents a push operation on stack.



where $\gamma_i = \gamma_i^1 \wedge \dots \wedge \gamma_i^{h_i}$ and $R_i = \{x_1, \dots, x_m\}$. The first and last micro-transitions, corresponding to the reset of a new clock ζ and checking of constraint $\zeta = 0$ ensure that all micro-transitions in the sequence occur simultaneously. We have a point in V_i for each micro-transition (excluding the ε -micro-transitions between $\delta_i^{x_j}$). Hence, V_i consists of a sequence $\ell_i \rightarrow \ell_i^1 \rightarrow \dots \rightarrow \ell_i^{h_i} \rightarrow p_i \rightarrow r_i^1 \rightarrow \dots \rightarrow r_i^{g_i} \rightarrow r_i$ where g_i is the number of timing constraints corresponding to clocks reset during transition i and checked afterwards. Thus, the reset-loop on a clock is fired as many times (0 or more) as a constraint is checked on this clock until its next reset. This ensures that the STCW remains *simple*. Similarly, h_i is the number of timing constraints checked in γ_i . We have $\lambda(p_i) = a_i$ and all other points are labelled ε . The set V_0 encodes the initial resets of clocks that will be checked before being reset. So we let $R_0 = X$ and V_0 is $\ell_0 \rightarrow r_0^1 \rightarrow \dots \rightarrow r_0^{g_0} \rightarrow r_0$.

- for each $I \in \mathcal{I}$, the relation for timing constraints can be partitioned as $\curvearrowright^I = \curvearrowright^{s \in I} \uplus \biguplus_{x \in X \cup \{\zeta\}} \curvearrowright^{x \in I}$ where
 - $\curvearrowright^{\zeta \in I} = \{(\ell_i, r_i) \mid 0 \leq i \leq n\}$ and $I = [0, 0]$.
 - We have $p_i \curvearrowright^{s \in I} p_j$ if $\text{op}_i = \downarrow_b$ is a push and $\text{op}_j = \uparrow_b^I$ is the matching pop (same number of pushes and pops in $\text{op}_{i+1} \dots \text{op}_{j-1}$).
 - for each $0 \leq i < j \leq n$ such that the t -th conjunct of γ_j is $x \in I$ and $x \in R_i$ and $x \notin R_k$ for $i < k < j$, we have $r_i^s \curvearrowright^{x \in I} \ell_j^t$ for some $1 \leq s \leq g_i$. Therefore, every point ℓ_i^t with $1 \leq t \leq h_i$ is the target of a timing constraint. Moreover, every reset point r_i^s for $1 \leq s \leq g_i$ should be the source of a timing constraint. That is, denoting $\curvearrowright^x = \bigcup_{I' \in \mathcal{I}} \curvearrowright^{x \in I'}$, we must have $r_i^s \in \text{dom}(\curvearrowright^x)$ for some $x \in R_i$. Also, for each i , the reset points $r_i^1, \dots, r_i^{g_i}$ are grouped by clocks (as suggested by the sequence of micro-transitions simulating δ_i): if $1 \leq s < u < t \leq g_i$ and $r_i^s, r_i^t \in \text{dom}(\curvearrowright^x)$ for some $x \in R_i$ then $r_i^u \in \text{dom}(\curvearrowright^x)$. Finally, for each clock, we require that the timing constraints are well-nested: for all $u \curvearrowright^x v$ and $u' \curvearrowright^x v'$, with $u, u' \in V_i$, if $u < u'$ then $u' < v' < v$.

We denote by $\text{STCW}(\mathcal{S})$ the set of simple TCWs generated by \mathcal{S} and define the language of \mathcal{S} as the set of *realizable* STCWs, i.e., $\mathcal{L}(\mathcal{S}) = \text{Real}(\text{STCW}(\mathcal{S}))$. Indeed, this is equivalent to defining the language as the set of timed words accepted by \mathcal{S} , according to a usual operational semantics [1].

The STCW semantics of timed automata (TA) can be obtained from the above discussion by just ignoring the stack components (using **nop** operations only). To illustrate these ideas, a simple example of a timed automaton and an STCW that is generated by it is shown in Figure 1.

We now identify some important properties satisfied by STCWs generated from a TPDA. Let $\mathcal{V} = (V, \rightarrow, (\curvearrowright^I)_{I \in \mathcal{I}}, \lambda)$ be a STCW. We say that \mathcal{V} is *well timed* w.r.t. a set of clocks Y and a stack s if for each interval $I \in \mathcal{I}$, the \curvearrowright^I relation can be partitioned as $\curvearrowright^I = \curvearrowright^{s \in I} \uplus \biguplus_{x \in Y} \curvearrowright^{x \in I}$ where

- (T₁) the stack relation $\curvearrowright^s = \bigcup_{I \in \mathcal{I}} \curvearrowright^{s \in I}$ corresponds to the matching push-pop events, hence it is well-nested: for all $i \curvearrowright^s j$ and $i' \curvearrowright^s j'$, if $i < i' < j$ then $i' < j' < j$.
- (T₂) For each clock $x \in Y$, the relation $\curvearrowright^x = \bigcup_{I \in \mathcal{I}} \curvearrowright^{x \in I}$ corresponds to the timing constraints for clock x and is well-nested: for all $i \curvearrowright^x j$ and $i' \curvearrowright^x j'$, if $i < i'$ are in the same x -reset block (i.e., a maximal consecutive sequence $i_1 < \dots < i_n$ of positions in the domain of \curvearrowright^x), and $i < i' < j$, then $i' < j' < j$. Each guard should be

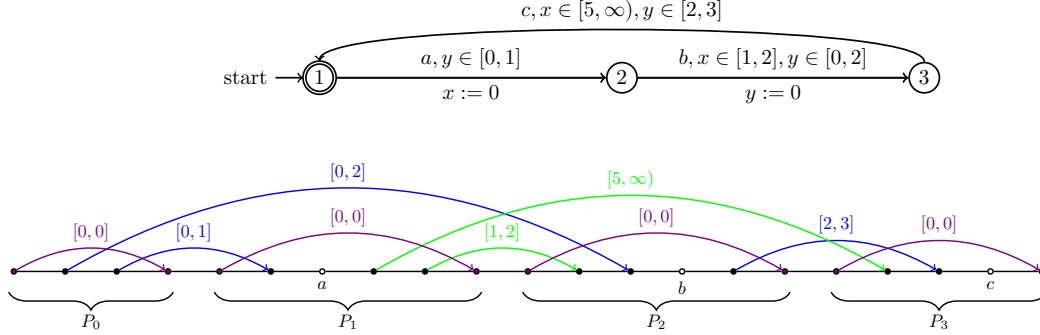


Figure 1: A timed automaton (top) with 2 clocks x, y . An STCW generated by an accepting run of the TA is depicted just below. The blue edges represent matching relations induced by clock y , while the green represent those induced by clock x . The violet edges are the $[0, 0]$ timing constraints for the extra clock ζ ensuring that all points in some V_i representing transition i occur precisely at the same time. Black lines are process edges.

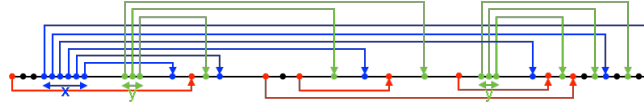


Figure 2: Well-timed: relations \sim^s in red, \sim^x in blue and \sim^y in green.

matched with the closest reset block on its left: for all $i \sim^x j$ and $i' \sim^x j'$, if $i < i'$ are not in the same x -reset block then $j < i'$ (see Figure 2).

It is then easy to check that STCWs defined by a TPDA with set of clocks X are well-timed for the set of clocks $Y = X \cup \{\zeta\}$, i.e., satisfy the properties above: (note that we obtain the same for TA by just ignoring the stack edges, i.e., (T_1) above)

Lemma 2.2. *Simple TCWs defined by a TPDA with set of clocks X are well-timed wrt. set of clock $Y = X \cup \{\zeta\}$, i.e., satisfy properties (T_1) and (T_2) .*

Proof. The first condition (T_1) is satisfied by $\text{STCW}(\mathcal{S})$ by definition. For (T_2) , let $i \sim^x j$ and $i' \sim^x j'$ for some clock $x \in X$. If i, i' are points in the same x -reset block for some $x \in X$, then by construction of $\text{STCW}(\mathcal{S})$, if $i < i'$ then $i' < j' < j$ which gives well nesting. Similarly, if $i < i'$ are points in different x -reset blocks, then by definition of $\text{STCW}(\mathcal{S})$, we have $j < i'$. Also, it is clear that the new clock ζ satisfies (T_2) . \square

3. BOUNDING THE WIDTH OF GRAPH BEHAVIORS OF TIMED SYSTEMS

In this section, we check if the graphs (STCWs) introduced in the previous section have a bounded tree-width. As a first step towards that, we introduce *special tree terms* (STTs) from Courcelle [8] and their semantics as labeled graphs.

3.1. Preliminaries: special tree terms and tree-width. It is known [8] that special tree terms using at most K colors (K -STTs) define graphs of “special” tree-width at most $K - 1$. Formally, a (Σ, Γ) -labeled graph is a tuple $G = (V_G, (E_\gamma)_{\gamma \in \Gamma}, \lambda)$ where $\lambda: V_G \rightarrow \Sigma$ is the vertex labeling and $E_\gamma \subseteq V_G^2$ is the set of edges for each label $\gamma \in \Gamma$. Special tree terms form an algebra to define labeled graphs. The syntax of K -STTs over (Σ, Γ) is given by

$$\tau ::= (i, a) \mid \text{Add}_{i,j}^\gamma \tau \mid \text{Forget}_i \tau \mid \text{Rename}_{i,j} \tau \mid \tau \oplus \tau$$

where $a \in \Sigma$, $\gamma \in \Gamma$ and $i, j \in [K] = \{1, \dots, K\}$ are colors. The semantics of each K -STT is a colored graph $\llbracket \tau \rrbracket = (G_\tau, \chi_\tau)$ where G_τ is a (Σ, Γ) -labeled graph and $\chi_\tau: [K] \rightarrow V_G$ is a partial injective function assigning a vertex of G_τ to some colors.

- $\llbracket (i, a) \rrbracket$ consists of a single a -labeled vertex with color i .
- $\text{Add}_{i,j}^\gamma$ adds a γ -labeled edge to the vertices colored i and j (if such vertices exist).
Formally, if $\llbracket \tau \rrbracket = (V_G, (E_\gamma)_{\gamma \in \Gamma}, \lambda, \chi)$ then $\llbracket \text{Add}_{i,j}^\alpha \tau \rrbracket = (V_G, (E'_\gamma)_{\gamma \in \Gamma}, \lambda, \chi)$ with $E'_\gamma = E_\gamma$ if $\gamma \neq \alpha$ and $E'_\alpha = \begin{cases} E_\alpha & \text{if } \{i, j\} \not\subseteq \text{dom}(\chi) \\ E_\alpha \cup \{(\chi(i), \chi(j))\} & \text{otherwise.} \end{cases}$
- Forget_i removes color i from the domain of the color map.
Formally, if $\llbracket \tau \rrbracket = (V_G, (E_\gamma)_{\gamma \in \Gamma}, \lambda, \chi)$ then $\llbracket \text{Forget}_i \tau \rrbracket = (V_G, (E_\gamma)_{\gamma \in \Gamma}, \lambda, \chi')$ with $\text{dom}(\chi') = \text{dom}(\chi) \setminus \{i\}$ and $\chi'(j) = \chi(j)$ for all $j \in \text{dom}(\chi')$.
- $\text{Rename}_{i,j}$ exchanges the colors i and j .
Formally, if $\llbracket \tau \rrbracket = (V_G, (E_\gamma)_{\gamma \in \Gamma}, \lambda, \chi)$ then $\llbracket \text{Rename}_{i,j} \tau \rrbracket = (V_G, (E_\gamma)_{\gamma \in \Gamma}, \lambda, \chi')$ with $\chi'(\ell) = \chi(\ell)$ if $\ell \in \text{dom}(\chi) \setminus \{i, j\}$, $\chi'(i) = \chi(j)$ if $j \in \text{dom}(\chi)$ and $\chi'(j) = \chi(i)$ if $i \in \text{dom}(\chi)$.
- Finally, \oplus constructs the disjoint union of the two graphs provided they use different colors. This operation is undefined otherwise.
Formally, if $\llbracket \tau_i \rrbracket = (G_i, \chi_i)$ for $i = 1, 2$ and $\text{dom}(\chi_1) \cap \text{dom}(\chi_2) = \emptyset$ then $\llbracket \tau_1 \oplus \tau_2 \rrbracket = (G_1 \uplus G_2, \chi_1 \uplus \chi_2)$. Otherwise, $\tau_1 \oplus \tau_2$ is not a valid STT.

The *special tree-width* of a graph G is defined as the least K such that $G = G_\tau$ for some $(K + 1)$ -STT τ . See [8] for more details and its relation to tree-width. For TCWs, we have successor edges and \curvearrowright^I -edges carrying timing constraints, so we take $\Gamma = \{\rightarrow\} \cup \{(x, y) \mid x \in \mathbb{N}, y \in \overline{\mathbb{N}}\}$ with $\overline{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$. In this paper, we will actually make use of STTs with the following restricted syntax, which are sufficient and make our proofs simpler:

$$\text{atomicSTT} ::= (1, a) \mid \text{Add}_{1,2}^{x,y}((1, a) \oplus (2, b))$$

$$\tau ::= \text{atomicSTT} \mid \text{Add}_{i,j}^\rightarrow \tau \mid \text{Forget}_i \tau \mid \text{Rename}_{i,j} \tau \mid \tau \oplus \tau$$

with $a, b \in \Sigma_\varepsilon$, $0 \leq x < M$, $0 \leq y < M$ or $y = +\infty$ for some $M \in \mathbb{N}$ and $i, j \in [2K] = \{1, \dots, 2K\}$. The terms defined by this grammar are called (K, M) -STTs. Here, timing constraints are added directly between leaves in atomic STTs which are then combined using disjoint unions and adding successor edges. For instance, consider the 4-STT given below

$$\tau = \text{Forget}_3 \text{Add}_{1,3}^\rightarrow \text{Forget}_2 \text{Add}_{2,4}^\rightarrow \text{Add}_{3,2}^\rightarrow (\text{Add}_{1,2}^{2,\infty}((1, a) \oplus (2, c)) \oplus \text{Add}_{3,4}^{1,3}((3, b) \oplus (4, d)))$$

where $\text{Add}_{i,j}^\gamma((i, \alpha) \oplus (j, \beta))$, $i, j \in \mathbb{N}$, $\alpha, \beta \in \Sigma_\varepsilon$ stands for $\text{Rename}_{1,i} \text{Rename}_{2,j} \text{Add}_{1,2}^\gamma((1, \alpha) \oplus (2, \beta))$. The term τ is depicted as a binary tree on the left of Figure 3 and its semantics $\llbracket \tau \rrbracket$ is the STCW given below where only endpoints labelled a and d are colored, as the other two colors were “forgotten” by τ . Abusing notation, we will also use $\llbracket \tau \rrbracket$ for the graph G_τ ignoring the coloring χ_t .

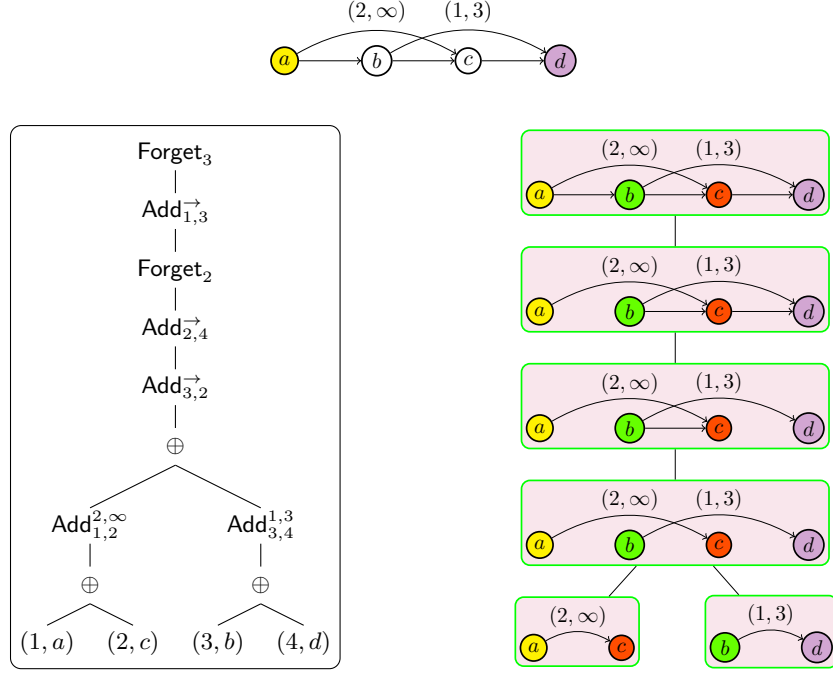


Figure 3: An STT and a simple TCW with its split-tree

3.2. Split-TCWs and split-game. There are many ways to show that a graph (in our case, a simple TCW) has a bounded special tree-width. We find it convenient to prove that a simple TCW has bounded special tree-width by playing a split-game, whose game positions are simple TCWs in which some successor edges have been cut, i.e., are missing. This approach has been used for graph behaviors generated by untimed pushdown systems in [11], but here we wish to apply it to reason about graph behaviors of timed systems. Formally, a *split-TCW* is a structure $(V, \rightarrow \cup \dashrightarrow, (\curvearrowright^I)_{I \in \mathcal{I}}, \lambda)$ where \rightarrow and \dashrightarrow are the present and absent successor edges (also called *holes*), respectively, such that $\rightarrow \cap \dashrightarrow = \emptyset$ and $(V, \rightarrow \cup \dashrightarrow, (\curvearrowright^I)_{I \in \mathcal{I}}, \lambda)$ is a TCW. Notice that, for a split-TCW, $\leq = \rightarrow \cup \dashrightarrow$ and $< = <^+$. A *block* or *factor* of a split-TCW is a maximal set of points of V connected by \rightarrow . We denote by $\text{EP}(\mathcal{V}) \subseteq V$ the set of left and right endpoints of blocks of \mathcal{V} . A left endpoint e is one for which there is no f with $f \rightarrow e$. Right endpoints are defined similarly. Points in $V \setminus \text{EP}(\mathcal{V})$ are called internal. If i is any point, $\text{Block}(i)$ denotes the unique block which contains i . The number of blocks is the *width* of \mathcal{V} : $\text{width}(\mathcal{V}) = 1 + |\dashrightarrow|$. TCWs may be identified with split-TCWs of width 1, i.e., with $\dashrightarrow = \emptyset$. A split-TCW is *atomic* if it consists of a single point ($|V| = 1$) or a single timing constraint with a hole ($V = \{p_1, p_2\}$, $p_1 \dashrightarrow p_2$, $p_1 \curvearrowright^I p_2$). In what follows, we will use the notation \curvearrowright for $\bigcup_{I \in \mathcal{I}} \curvearrowright^I$ when convenient.

The *split-game* is a two player turn based game $\mathcal{G} = (V_{\exists} \uplus V_{\forall}, E)$ where Eve's set of game positions V_{\exists} consists of all connected (wrt. $\rightarrow \cup \curvearrowright$) split-TCWs and Adam's set of game positions V_{\forall} consists of non-connected split-TCWs. The edges E of \mathcal{G} reflect the moves of the players. Eve's moves consist of splitting a block in two, i.e., removing one successor

edge in the graph. Adam's moves amount to choosing a connected component of the split-TCW. Atomic split-TCWs are terminal positions in the game: neither Eve nor Adam can move from an atomic split-TCW. A play on a split-TCW \mathcal{V} is a path in \mathcal{G} starting from \mathcal{V} and leading to an atomic split-TCW. The cost of the play is the maximum width of any split-TCW encountered in the path. Eve's objective is to minimize the cost, while Adam's objective is to maximize it.

A strategy for Eve from a split-TCW \mathcal{V} can be described with a *split-tree* T which is a binary tree labeled with split-TCWs satisfying:

- (1) The root is labeled by $\mathcal{V} = \text{labroot}(T)$.
- (2) Leaves are labeled by atomic split-TCWs.
- (3) Eve's move: Each unary node is labeled with some connected (wrt. $\rightarrow \cup \curvearrowright$) split-TCW \mathcal{V} and its child is labeled with some \mathcal{V}' obtained by splitting a block of \mathcal{V} in two, i.e., by removing one successor edge. Thus, $\text{width}(\mathcal{V}') = 1 + \text{width}(\mathcal{V})$.
- (4) Adam's move: Each binary node is labeled with some non connected split-TCW $\mathcal{V} = \mathcal{V}_1 \uplus \mathcal{V}_2$ where \mathcal{V}_1 and \mathcal{V}_2 are the labels of its children. Note that $\text{width}(\mathcal{V}) = \text{width}(\mathcal{V}_1) + \text{width}(\mathcal{V}_2)$.

The *width* of a split-tree T , denoted $\text{width}(T)$, is the maximum width of the split-TCWs labeling the nodes of T . In other words, the cost of the strategy encoded by T is $\text{width}(T)$ and this is the maximal cost of the plays starting from \mathcal{V} and following this strategy. A K -split-tree is a split-tree of width at most K .

The *split-width* of a simple (split-)TCW \mathcal{V} is the minimal cost of Eve's (positional) strategies starting from \mathcal{V} . In other words it is the minimum width of any split-tree for the simple TCW. Notice that Eve has a strategy to decompose a TCW \mathcal{V} into *atomic* split-TCWs if and only if \mathcal{V} is *simple*, i.e., at most one timing constraint is attached to each point. An example of a split-tree is given in Figure 3 (right). Observe that the width of the split-tree is 4. Hence the split-width of the simple TCW labeling the root is at most four.

Let STCW^K (resp. $\text{STCW}^{K,M}$) denote the set of simple TCWs with split-width bounded by K (resp. and using constants at most M) over the fixed alphabet Σ . The crucial link between special tree-width and split-width is given by the following lemma.

Lemma 3.1. *A (split) STCW of split-width at most K has special tree-width at most $2K - 1$.*

Intuitively, we only need to keep colors for end-points of blocks. Hence, each block of an STCW \mathcal{V} needs at most two colors and if the width of \mathcal{V} is at most K then we need at most $2K$ colors. From this it can be shown that a strategy of Eve of cost at most K can be encoded by a $2K$ -STT, which gives a special tree-width of at most $2K - 1$.

Proof. Let $\mathcal{V} = (V, \rightarrow \cup \dashrightarrow, (\curvearrowright^I)_{I \in \mathcal{I}}, \lambda)$ be a split-TCW. Recall that we denote by $\text{EP}(\mathcal{V})$ the subset of events that are endpoints of blocks of \mathcal{V} . A left endpoint is an event $e \in V$ such that there are no f with $f \rightarrow e$. We define similarly right endpoints. Note that an event may be both a left and right endpoint. The number of endpoints is at most twice the number of blocks: $|\text{EP}(\mathcal{V})| \leq 2 \cdot \text{width}(\mathcal{V})$.

We associate with every split-tree T of width at most K a $2K$ -STT \bar{T} such that $\llbracket \bar{T} \rrbracket = (\mathcal{V}, \chi)$ where $\mathcal{V} = \text{labroot}(T)$ is the label of the root of T and the range of χ is the set of endpoints of \mathcal{V} : $\text{Im}(\chi) = \text{EP}(\mathcal{V})$. Notice that $\text{dom}(\chi) \subseteq [2K]$ since \bar{T} is a $2K$ -STT. The construction is by induction on T .

Assume that $\text{labroot}(T)$ is atomic. Then it is either an internal event labeled $a \in \Sigma$, and we let $\bar{T} = (1, a)$. Or, it is a pair of events $e \curvearrowright^I f$ with a timing constraint $I = [c, d]$ and we let $\bar{T} = \text{Add}_{1,2}^{c,d}((1, \lambda(e)) \oplus (2, \lambda(f)))$.

If the root of T is a binary node and the left and right subtrees are T_1 and T_2 then $\text{labroot}(T) = \text{labroot}(T_1) \uplus \text{labroot}(T_2)$. By induction, for $i = 1, 2$ the STT \bar{T}_i is already defined and we have $\llbracket \bar{T}_i \rrbracket = (\text{labroot}(T_i), \chi_i)$. We first rename colors that are active in both STTs. To this end, we choose an injective map $f: \text{dom}(\chi_1) \cap \text{dom}(\chi_2) \rightarrow [2K] \setminus (\text{dom}(\chi_1) \cup \text{dom}(\chi_2))$. This is possible since $|\text{dom}(\chi_i)| = |\text{Im}(\chi_i)| = |\text{EP}(\text{labroot}(T_i))|$. Hence, $|\text{dom}(\chi_1)| + |\text{dom}(\chi_2)| = |\text{EP}(\text{labroot}(T))| \leq 2K$.

Assuming that $\text{dom}(f) = \{i_1, \dots, i_m\}$, we define

$$\bar{T} = \bar{T}_1 \oplus \text{Rename}_{i_1, f(i_1)} \cdots \text{Rename}_{i_m, f(i_m)} \bar{T}_2.$$

Finally, assume that the root of T is a unary node with subtree T' . Then, $\text{labroot}(T')$ is obtained from $\text{labroot}(T)$ by splitting one block, i.e., removing one word edge, say $e \rightarrow f$. We deduce that e and f are endpoints of $\text{labroot}(T')$, respectively right and left endpoints. By induction, the STT \bar{T}' is already defined. We have $\llbracket \bar{T}' \rrbracket = (\text{labroot}(T'), \chi')$ and $e, f \in \text{Im}(\chi')$. So let i, j be such that $\chi'(i) = e$ and $\chi'(j) = f$. We add the process edge with $\tau = \text{Add}_{i,j}^{\rightarrow} \bar{T}'$. Then we forget color i if e is no more an endpoint, and we forget j if f is no more an endpoint:

$$\tau' = \begin{cases} \tau & \text{if } e \text{ is still an endpoint,} \\ \text{Forget}_i \tau & \text{otherwise} \end{cases} \quad \bar{T} = \begin{cases} \tau' & \text{if } f \text{ is still an endpoint,} \\ \text{Forget}_j \tau' & \text{otherwise.} \end{cases}$$

□

3.3. Split-width for timed systems. Viewing terms as trees, our goal in the next section will be to construct tree automata to recognize sets of (K, M) -STTs, and thus capture the $(K \text{ split-width})$ bounded behaviors of a given system. A possible way to show that these capture *all* behaviors of the given system, is to show that we can find a K such that all the (graph) behaviors of the given system have a K -bounded split-width.

We do this now for a TPDA and also mention how to modify the proof for a timed automaton. In Section 6, we also show how it extends to multi-pushdown systems. In fact, we prove a slightly more general result, by showing that all well-timed split-STCWs defined in Section 2.3 for the set of clocks $Y = X \cup \{\zeta\}$ have bounded split-width (lifting the definition of well-timed to split-STCWs). From Lemma 2.2, it follows that the STCWs defined by a TPDA with set of clocks X are well-timed for the set of clocks $Y = X \cup \{\zeta\}$ and hence we obtain a bound on the split-width as required. Let $\text{STCW}(\mathcal{S})$ represent the STCWs that are defined by a TPDA \mathcal{S} .

Lemma 3.2. *The split-width of a well-timed STCW is bounded by $4|Y| + 2$.*

Proof. This lemma is proved by playing the split-width game between *Adam* and *Eve*. *Eve* should have a strategy to disconnect the word without introducing more than $4|Y| + 2$ blocks. The strategy of *Eve* uses three operations processsing the word from right to left.

Removing an internal point. If the last/right-most event on the word (say event j) is not the target of a \curvearrowright relation, then she will split the \rightarrow -edge before the last point, i.e., the edge between point j and its predecessor.

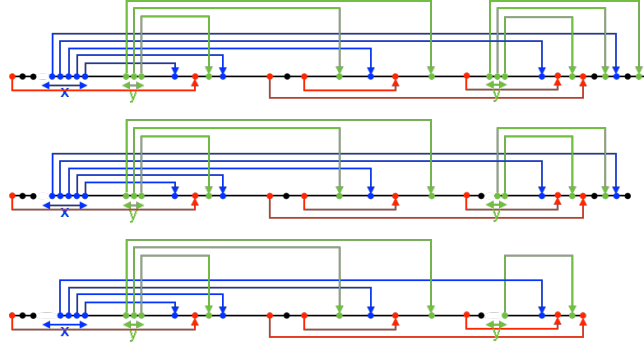


Figure 4: Removing timing constraints. At most one reset hole per clock. Edges below are stack edges. Clock edges labeled with x, y .

Removing a clock constraint. Assume that we have a timing constraint $i \curvearrowright^x j$ where j is the last point of the split-TCW. Then, by (T_2) we deduce that i is the *first* point of the *last reset block* for clock x . *Eve* splits three word-edges to detach the matching pair $i \curvearrowright^x j$: these three edges are those connected to i and j . Since the matching pair $i \curvearrowright^x j$ is atomic, *Adam* should continue the game from the remaining split-TCW \mathcal{V}' . Notice that we have now a hole instead of position i . We call this a reset-hole for clock x . For instance, starting from the split-TCW of Figure 2, and removing the last timing constraint of clock x , we get the split-TCW on top of Figure 4. Notice the reset hole at the beginning of the x reset block.

During the inductive process, we may have at most one such reset hole for each clock $x \in Y$. Note that the first time we remove the last point which is a timing constraint for clock x , we create a hole in the last reset block of x which contains a sequence of reset points for x , by removing two edges. This hole is created by removing the leftmost point in the reset block. As we keep removing points from the right which are timing constraints for x , this hole widens in the reset block, by removing each time, just one edge in the reset block. Continuing the example, if we detach the last internal point and then the timing constraint of clock y we get the split-TCW in the middle of Figure 4. Now, we have one reset hole for clock x and one reset hole for clock y .

We continue by removing from the right, one internal point, one timing constraint for clock x , one timing constraint for clock y , and another internal point, we get the split-TCW at the bottom of Figure 4. Notice that we still have a *single* reset hole for each clock x, y .

Removing a push-pop edge. Assume now that the last event is a pop: we have $i \curvearrowright^s j$ where j is the last point of the split-TCW. If there is already a hole before the push event i or if i is the first point of the split-TCW, then we split after i and before j to detach the atomic matching pair $i \curvearrowright^s j$. *Adam* should choose the remaining split-TCW and the game continues.

Note that if i is not the first point of the split-TCW and there is no hole before i , we cannot proceed as we did in the case of clock constraints, since this would create a push-hole and the pushes are not arranged in blocks as the resets. Hence, removing push-pop edges as we removed timing constraints would create an unbounded number of holes. Instead, we split the TCW just before the matching push event i . Since push-pop edges are well-nested (T_1) and since j is the last point of the split-TCW, there are no push-pop edges crossing

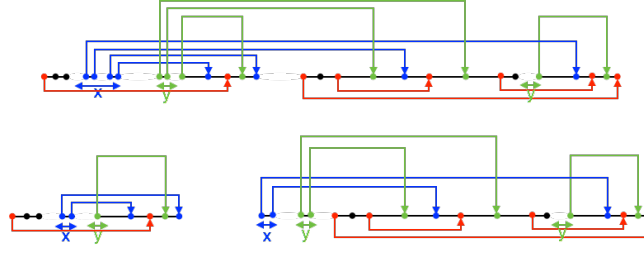
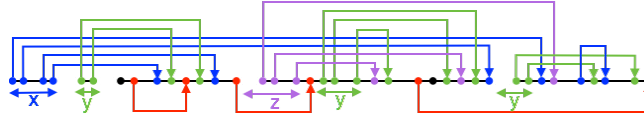


Figure 5: Splitting the TCW at a position with no crossing stack edge.

Figure 6: A split-TCW satisfying (T_1-T_2) and (I_1-I_2) .

position i : $i' \curvearrowright^s j'$ and $i' < i$ implies $j' < i$. Hence, only clock constraints may cross position i .

Consider some clock x having timing constraints crossing position i . All these timing constraints come from the last reset block B_x of clock x which is before position i . Moreover, these resets form the left part of the reset block B_x . We detach this left part with two splits, one before the reset block B_x and one after the last reset of block B_x whose timing constraint crosses position i . We proceed similarly for each clock of Y . Recall that we also split the TCW just before the push event i . As a result, the TCW is not connected anymore. Notice that we have used at most $2|Y| + 1$ new splits to disconnect the TCW. For instance, from the bottom split-TCW of Figure 4, applying the procedure above, we obtain the split-TCW on top of Figure 5 which has two connected components. Notice that to detach the left part of the reset block of clock x we only used one split since there was already a reset hole at the beginning of this block. The two connected components are depicted separately at the bottom of Figure 5.

Invariant. The split-TCW at the bottom right of Figure 5 is representative of the split-TCW which may occur during the split-game using the strategy of *Eve* described above. These split-TCW satisfy the following invariant.

- (I_1) The split-TCW starts with reset blocks, at most one for each clock in Y . For instance, the split-TCW of Figure 6 starts with two *reset blocks*, one for clock x and one for clock y (see the two hanging reset blocks on the left, one of x and one of y).
- (I_2) After these reset blocks, the split-TCW may have reset holes, at most one for each clock in Y . For instance, the split-TCW of Figure 6 has two *reset holes*, one for clock z and one for clock y .

A reset hole for clock x is followed by the last reset block of clock x , if any. Hence, for all timing constraints $i \curvearrowright^x j$ such that j is on the right of the hole, the reset event i is in the reset block that starts just after the hole.

Claim 3.3. *A split-TCW satisfying (I_1-I_2) has at most $2|Y| + 1$ blocks. It is disconnected by Eve's strategy using at most $2|Y| + 1$ new splits, and the resulting connected components satisfy (I_1-I_2) . Therefore, its split-width is at most $4|Y| + 2$.*

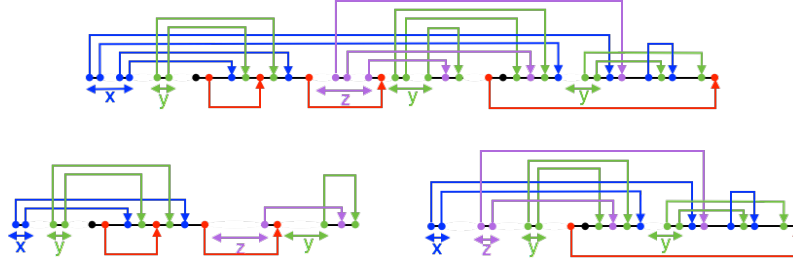


Figure 7: Splitting the TCW at a position with no crossing stack edge.

Proof. Let us check that starting from a split-TCW satisfying (I_1-I_2) , we can apply *Eve's* strategy and the resulting connected components also satisfy (I_1-I_2) . This is trivial when the last point is internal in which case *Eve* makes one split before this last point.

In the second case, the last event checks a timing constraint for some clock x : we have $i \curvearrowright^x j$ and j is the last event. If there is already a reset hole for clock x in the split-TCW, then by (I_2) and (T_2) , the reset event i must be just after the reset hole for clock x . So with at most two new splits *Eve* detaches the atomic edge $i \curvearrowright^x j$ and the resulting split-TCW satisfies the invariants. If there is no reset hole for clock x then we consider the last reset block B_x for clock x . By (T_2) , i must be the first event of this block. Either B_x is one of the first reset blocks of the split-TCW (I_1) and *Eve* detaches with at most two splits the atomic edge $i \curvearrowright^x j$. Or *Eve* detaches this atomic edge with at most three splits, creating a reset hole for clock x in the resulting split-TCW. In both cases, the resulting split-TCW satisfies the invariants.

The third case is when the last event is a pop event: $i \curvearrowright^s j$ and j is the last event. Then there are two subcases: either, there is a hole before event i then *Eve* detaches with two splits the atomic edge $i \curvearrowright^s j$ and the resulting split-TCW satisfies the invariants. Or we cut the split-TCW before position i . Notice that no push-pop edges cross i : if $i' \curvearrowright^s j'$ and $i' < i$ then $j' < i$. As above, for each clock x having timing constraints crossing position i , we consider the last reset block B_x for clock x which is before position i . The resets of the timing constraints for clock x crossing position i form a left factor of the reset block B_x . We detach this left factor with at most two splits. We proceed similarly for each clock of Y . The resulting split-TCW is not connected anymore and we have used at most $2|Y| + 1$ more splits. For instance, if we split the TCW of Figure 6 just before the last push following the procedure described above, we get the split-TCW on top of Figure 7. This split-TCW is not connected and its left and right connected components are drawn below.

To see that the invariants are maintained by the connected components, let us inspect the splitting of block B_x . First, B_x could be one of the beginning reset blocks (I_1) . This is the case for clock x in Figures 6 and 7. In which case *Eve* use only one split to divide B_x in B_x^1 and B_x^2 . The left factor B_x^1 corresponds to the edges crossing position i and will form one of the reset block (I_1) of the right connected component. On the other hand, the suffix B_x^2 stays a reset block of the left connected component. Second, B_x could follow a reset hole for clock x . This is the case for clock z in Figures 6 and 7. In which case again *Eve* only needs one split to detach the left factor B_x^1 which becomes a reset block (I_1) of the right component. The reset hole before B_x stays in the left component. Finally, assume that B_x is neither a beginning reset block (I_1) , nor follows a reset hole for clock x . This is

the case for clock y in Figures 6 and 7. Then *Eve* detaches the left factor B_x^1 which becomes a reset block (l_1) of the right component and creates a reset hole in the left component. \square

This claim along with the strategy ends the proof of the lemma. \square

Now, if the STCW is from a timed automaton then, \curvearrowright^s is empty and *Eve*'s strategy only has the first two cases above. Thus, we obtain a bound of $|Y| + 3$ on split-width for timed automata. Thus, we have our second main contribution of this paper and the main theorem of this section, namely, Theorem 3.4 (2).

Theorem 3.4. *Given a timed system \mathcal{S} using a set of clocks X , all words in its STCW language have split-width bounded by K , i.e., $\text{STCW}(\mathcal{S}) \subseteq \text{STCW}^K$, where*

- (1) $K = |X| + 4$ if \mathcal{S} is a timed automaton,
- (2) $K = 4|X| + 6$ if \mathcal{S} is a timed pushdown automaton,

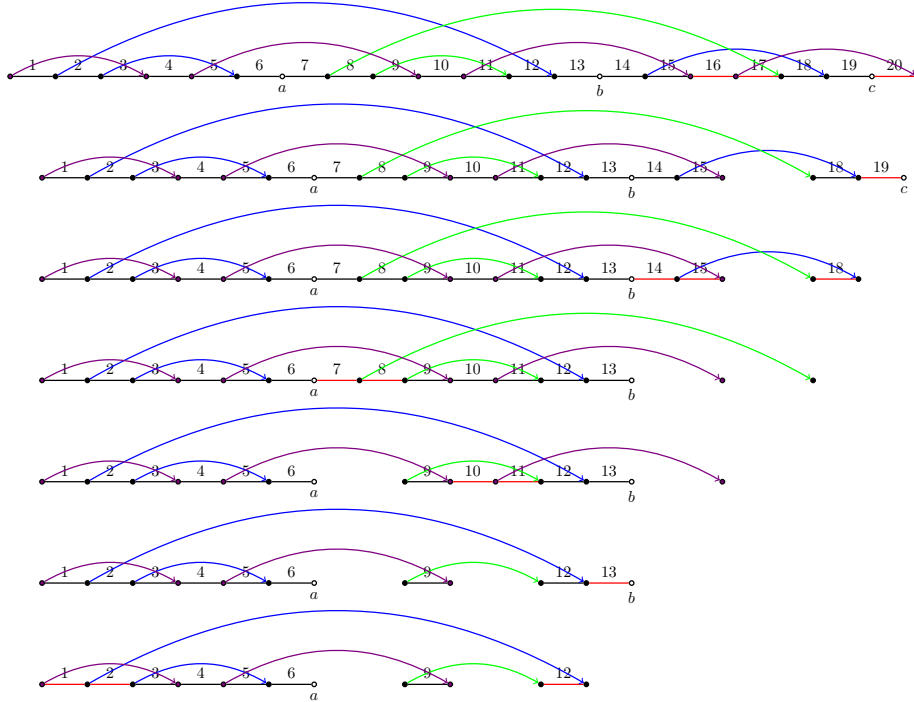


Figure 8: The first 4 steps of the split-game on the STCW of Figure 1. Note that we do not write the time intervals or transition names as they are irrelevant for this game. In the first step, *Eve* detaches the last timing constraint by cutting the 3 successor edges 16, 17 and 20. Adam chooses the non atomic STCW of the second line. In the second step, *Eve* detaches the last internal event labelled c by cutting edge 19. The resulting STCW is on the third line, from which *Eve* detaches the last timing constraint by cutting edges 14, 15, 18. She continues by cutting edges 7, 8 to detach the green timing constraint, and so on. Notice that we have three clocks $X = \{x, y, \zeta\}$ and that the maximal number of blocks $|X| + 3 = 6$ is reached on the last STCW when *Eve* cuts edges 1, 2 and 12.

Figure 8 illustrates our strategy starting from the STCW depicted in Figure 1, generated by the TA shown in that figure.

4. TREE AUTOMATA FOR VALIDITY

Fix $K, M \geq 2$. Not all graphs defined by (K, M) -STTs are realizable TCWs. Indeed, if τ is such an STT, the edge relation \rightarrow may have cycles or may be branching, which is not possible in a TCW. Also, the timing constraints given by \curvearrowright need not comply with the \rightarrow relation: for instance, we may have a timing constraint $e \curvearrowright f$ with $f \rightarrow^+ e$. Moreover, some terms may define graphs denoting TCWs which are not realizable. So we construct the tree automaton $\mathcal{A}_{\text{valid}}^{K,M}$ to check for validity. We will see at the end of the section (Corollary 4.6) that we can restrict $\mathcal{A}_{\text{valid}}^{K,M}$ further so that it accepts terms denoting *simple* TCWs of split-width $\leq K$.

First we show that, since we have only closed intervals in the timing constraints, considering integer timestamps is sufficient for realizability.

Lemma 4.1. *Let $W = (V, \rightarrow, (\curvearrowright^I)_{I \in \mathcal{I}(M)}, \lambda)$ be a TCW using only closed intervals in its timing constraints. Then, W is realizable iff there exists an integer valued timestamp map satisfying all timing constraints.*

Proof. Consider two non-negative real numbers $a, b \in \mathbb{R}_+$ and let $i = \lfloor a \rfloor$ and $j = \lfloor b \rfloor$ be their integral parts. Then, $j - i - 1 < b - a < j - i + 1$. It follows that for all closed intervals I with integer bounds, we have $b - a \in I$ implies $j - i \in I$.

Assume there exists a non-negative real-valued timestamp map $\text{ts}: V \rightarrow \mathbb{R}_+$ satisfying all timing constraints of W . From the remark above, we deduce that $\lfloor \text{ts} \rfloor: V \rightarrow \mathbb{N}$ also realizes all timing constraints of W . The converse direction is clear. \square

Consider a set of colors $P \subseteq \{1, \dots, K\}$. For each $i \in P$ we let $i^+ = \min\{j \in P \cup \{\infty\} \mid i < j\}$ and $i^- = \max\{j \in P \cup \{0\} \mid j < i\}$. If P is not clear from the context, then we write $\text{next}_P(i)$ and $\text{prev}_P(i)$.

We say that a (K, M) -STT is *monotonic* if for every subterm of the form $\text{Add}_{i,j}^{\rightarrow} \tau'$ (resp. $\text{Add}_{i,j}^{\curvearrowright^I} \tau'$ or $\text{Rename}_{i,j} \tau'$) we have $j = \text{next}_P(i)$ (resp. $i < j$ or $\text{prev}_P(i) < j < \text{next}_P(i)$), where P is the set of active colors in τ' .

In the following, we prove one of our main results viz., the construction of the tree automaton $\mathcal{A}_{\text{valid}}^{K,M}$ for checking the validity of the STTs.

Theorem 4.2. *We can build a tree automaton $\mathcal{A}_{\text{valid}}^{K,M}$ with $M^{\mathcal{O}(K)}$ states such that $\mathcal{L}(\mathcal{A}_{\text{valid}}^{K,M})$ is the set of monotonic (K, M) -STTs τ such that $\llbracket \tau \rrbracket$ is a realizable TCW and the endpoints of $\llbracket \tau \rrbracket$ are the only colored points.*

A (K, M) -STT τ defines a colored graph $\llbracket \tau \rrbracket = (G_\tau, \chi_\tau)$ where the graph G_τ is written $G_\tau = (V, \rightarrow, (\curvearrowright^I)_{I \in \mathcal{I}(M)}, \lambda)$. Examples of STTs and their semantics are given in Table 1. The graph $\llbracket \tau \rrbracket$ (or more precisely G_τ) is a TCW if it satisfies several MSO-definable conditions. First, \rightarrow should be the successor relation of a total order on V . Second, each timing constraint should be compatible with the total order: $\curvearrowright^I \subseteq \rightarrow^+$. Also, the TCW $\llbracket \tau \rrbracket$ should be realizable. These graph properties are MSO definable. Since a graph $\llbracket \tau \rrbracket$ has an MSO-interpretation in the tree τ , we deduce that there is a tree automaton that accepts all (K, M) -STTs τ such that $\llbracket \tau \rrbracket$ is a realizable TCW (cf. [9]).

τ_1	τ_2	τ_3
$\begin{array}{c} \text{Add}_{1,6}^{\curvearrowright[2,\infty]} \\ \\ \oplus \\ / \quad \backslash \\ (1, a) \quad (6, b) \end{array}$	$\begin{array}{c} \text{Forget}_4 \\ \\ \text{Add}_{3,4}^{\rightarrow} \\ \\ \oplus \\ / \quad \backslash \\ \text{Add}_{2,3}^{\curvearrowright[1,3]} \quad \text{Add}_{4,5}^{\rightarrow} \\ \quad \quad \\ \oplus \quad \quad \text{Add}_{4,5}^{\curvearrowright[2,3]} \\ / \quad \backslash \quad \quad \\ (2, c) \quad (3, d) \quad \oplus \\ \quad \quad \quad (4, c) \quad (5, d) \end{array}$	$\begin{array}{c} \oplus \\ / \quad \backslash \\ \text{Rename}_{3,5} \quad \text{Add}_{3,4}^{\rightarrow} \\ \quad \quad \\ \text{Add}_{1,2}^{\rightarrow} \quad \text{Add}_{3,4}^{\curvearrowright[2,\infty]} \\ \quad \quad \\ \text{Forget}_5 \quad \oplus \\ \quad \quad / \quad \backslash \\ \text{Add}_{5,6}^{\rightarrow} \quad (3, a) \quad (4, b) \\ \\ \oplus \\ / \quad \backslash \\ \tau_1 \quad \tau_2 \end{array}$
$\begin{array}{cc} P & 1 \quad 6 \\ \text{tsm} & 0 \quad 2 \end{array}$	$\begin{array}{ccccccc} P & 2 & 3 & \longrightarrow & 5 \\ \text{tsm} & 2 & \text{ac} & 1 & \text{ac} & 0 \end{array}$	$\begin{array}{ccccccc} \longrightarrow & 2 & 3 & \longrightarrow & 4 & 5 & \longrightarrow & 6 \\ \text{uc} & 2 & \text{ac} & 3 & \text{ac} & 1 & \text{ac} & 1 \end{array}$

Table 1: The second line gives the tree representations of three monotonic $(6, 4)$ -STTs τ_1 , τ_2 , τ_3 , the third line gives their semantics $\llbracket \tau \rrbracket = (G_\tau, \chi_\tau)$ together with a realization ts , the fourth line gives possible states q of $\mathcal{A}_{\text{valid}}^{K,M}$ after reading the terms. The boolean maps sd and ac of a state are represented as follows: $\text{sd}(i) = \text{tt}$ iff there is a solid edge from i to i^+ , $\text{ac}(i) = \text{tt}$ iff the tag “ac” is between $\text{tsm}(i)$ and $\text{tsm}(i^+)$.

But this is not exactly what we want/need. So we provide in the proof below a direct construction of a tree automaton checking validity. There are several reasons for directly constructing the tree automaton $\mathcal{A}_{\text{valid}}^{K,M}$. First, this allows to have a clear upper-bound on the number of states of $\mathcal{A}_{\text{valid}}^{K,M}$ (this would be quite technical via MSO). Second, simplicity of the TCW and the bound on split-width can be enforced with no additional cost.

Proof. A state of $\mathcal{A}_{\text{valid}}^{K,M}$ will be an abstraction of the graph defined by the STT read so far. The finite abstraction will keep only the colored points of the graph. We will only accept *monotonic* terms for which the natural order on the active colors coincides with the order of the corresponding vertices in the final TCW. The monotonicity ensures that the graph defined by the STT is in fact a split-TCW.

Moreover, to ensure realizability of the TCW defined by a term, we will guess timestamps of vertices modulo M . We also guess while reading a subterm whether the time elapsed between two consecutive active colors is *big* ($\geq M$) or *small* ($< M$). Then, the automaton has to check that all these guesses are coherent and using these values it will check that every timing constraint is satisfied.

Formally, states of $\mathcal{A}_{\text{valid}}^{K,M}$ are tuples of the form $q = (P, \text{sd}, \text{tsm}, \text{ac})$, where $P \subseteq \{1, \dots, K\}$, $\text{sd}, \text{ac}: P \rightarrow \mathbb{B}$ and $\text{tsm}: P \rightarrow [M] = \{0, \dots, M-1\}$. The number of states of the tree automaton thus depends on P , and hence is $M^{\mathcal{O}(K)}$.

Examples are given in Table 1. The intuition is as follows. When reading bottom-up a (K, M) -STT τ with $\llbracket \tau \rrbracket = (V, \rightarrow, \lambda, (\curvearrowright^I)_{I \in \mathcal{I}(M)}, \chi)$, the automaton $\mathcal{A}_{\text{valid}}^{K,M}$ will reach a state $q = (P, \text{sd}, \text{ts}, \text{ac})$ such that

- (A₁) $P = \text{dom}(\chi)$ is the set of *active* (non forgotten) colors in τ and $\text{EP}(\llbracket \tau \rrbracket) \subseteq \chi(P)$: the endpoints of $\llbracket \tau \rrbracket$ are colored.
- (A₂) For all $i \in P$, we have $\text{sd}(i) = \text{tt}$ iff $i^+ \neq \infty$ and $\chi(i) \rightarrow^+ \chi(i^+)$ in $\llbracket \tau \rrbracket$.
- (A₃) Let $\dashrightarrow = \{(\chi(i), \chi(i^+)) \mid i \in P \wedge i^+ \neq \infty \wedge \text{sd}(i) = \text{ff}\}$. Then, $(\llbracket \tau \rrbracket, \dashrightarrow)$ is a split-TCW, i.e., $< = (\rightarrow \cup \dashrightarrow)^+$ is a total order on V , timing constraints in $\llbracket \tau \rrbracket$ are $<$ -compatible $\curvearrowright^I \subseteq <$ for all I , the *direct successor* relation of $<$ is $\leq = \rightarrow \cup \dashrightarrow$ and $\rightarrow \cap \dashrightarrow = \emptyset$.
- (A₄) There exists a timestamp map $\text{ts}: V \rightarrow \mathbb{N}$ such that
 - all constraints are satisfied: $\text{ts}(v) - \text{ts}(u) \in I$ for all $u \curvearrowright^I v$
 - time is non-decreasing: $\text{ts}(u) \leq \text{ts}(v)$ for all $u \leq v$
 - (tsm, ac) is the modulo M abstraction of ts : for all $i \in P$ we have $\text{tsm}(i) = \text{ts}(\chi(i))[M]$ and $\text{ac}(i) = \text{tt}$ iff $i^+ \neq \infty$ and $\text{ts}(\chi(i^+)) - \text{ts}(\chi(i)) = (\text{tsm}(i^+) - \text{tsm}(i))[M] < M$.

We say that q is a *realizable abstraction* of τ if it satisfies conditions (A₁–A₄).

Indeed, the finite state automaton $\mathcal{A}_{\text{valid}}^{K,M}$ cannot store the timestamp map ts witnessing realizability. Instead, it stores the modulo M abstraction (tsm, ac) . We will see that $\mathcal{A}_{\text{valid}}^{K,M}$ can check realizability based on the abstraction (tsm, ac) of ts and can maintain this abstraction while reading the term bottom-up.

We introduce some notations. Let $q = (P, \text{sd}, \text{tsm}, \text{ac})$ be a state and let $i, j \in P$ with $i \leq j$. We define $d(i, j) = (\text{tsm}(j) - \text{tsm}(i))[M]$ and $D(i, j) = \sum_{k \in P \mid i \leq k < j} d(k, k^+)$. We also define $\text{AC}(i, j) = \bigvee_{k \in P \mid i \leq k < j} \text{ac}(k)$. If the state is not clear from the context, then we write $d_q(i, j)$, $D_q(i, j)$, $\text{AC}_q(i, j)$. For instance, with the state q_3 corresponding to the term τ_3 of Table 1, we have $\text{AC}(1, 5) = \text{tt}$, $d(1, 5) = 1$ and $D(1, 5) = 5 = \text{ts}(5) - \text{ts}(1)$ is the *accurate* value of the time elapsed. Whereas, $\text{AC}(2, 6) = \text{ff}$ and $d(2, 6) = 0$, $D(2, 6) = 4$ are both *strict modulo- M under-approximations* of the time elapsed $\text{ts}(6) - \text{ts}(2) = 8$.

Claim 4.3. *Let q be a state and τ be an STT. Then the following hold:*

- (1) *Assume (A₁–A₃) are satisfied. Then, for all $i, j \in P$ we have $i < j$ iff $\chi(i) < \chi(j)$: the natural ordering on colors coincide with the ordering of colored points in the split-TCW $(\llbracket \tau \rrbracket, \dashrightarrow)$.*
- (2) *Assume that ts is a timestamp map satisfying items 2 and 3 of (A₄). Then, for all $i, j \in P$ such that $i \leq j$, we have $d(i, j) = D(i, j)[M] = (\text{ts}(\chi(j)) - \text{ts}(\chi(i)))[M]$ and $d(i, j) \leq D(i, j) \leq \text{ts}(\chi(j)) - \text{ts}(\chi(i))$ (d and D give modulo M under-approximations of the actual time elapsed). Moreover, AC tells whether D gives the accurate elapse of time:*

$$\begin{aligned} \text{ac}(i) = \text{tt} &\iff d(i, i^+) = \text{ts}(\chi(i^+)) - \text{ts}(\chi(i)) \\ \text{AC}(i, j) = \text{tt} &\iff D(i, j) = \text{ts}(\chi(j)) - \text{ts}(\chi(i)) \\ \text{AC}(i, j) = \text{ff} &\implies \text{ts}(\chi(j)) - \text{ts}(\chi(i)) \geq M \end{aligned}$$

Proof. 1. From (A₂–A₃) we immediately get $\chi(i) < \chi(i^+)$ for all $i \in P$ such that $i^+ \neq \infty$. By transitivity we obtain $\chi(i) < \chi(j)$ for all $i, j \in P$ with $i < j$. Since $<$ is a strict total order on V , we deduce that, if $\chi(i) < \chi(j)$ for some $i, j \in P$, then $j \leq i$ is not possible.

2. Let $i, j \in P$ with $i \leq j$. Using items 2 and 3 of (A_4) we get

$$d(i, j) = (\text{tsm}(j) - \text{tsm}(i))[M] = (\text{ts}(\chi(j))[M] - \text{ts}(\chi(i))[M])[M] = (\text{ts}(\chi(j)) - \text{ts}(\chi(i)))[M].$$

Applying this equality for every pair (k, k^+) such that $i \leq k < j$ we get $D(i, j)[M] = (\text{ts}(\chi(j)) - \text{ts}(\chi(i)))[M]$. Since ts is non-decreasing (item 2 of A_4), it follows that $d(i, j) \leq D(i, j) \leq \text{ts}(\chi(j)) - \text{ts}(\chi(i))$.

Now, using again (A_4) we obtain $\text{ac}(k) = \text{tt}$ iff $d(k, k^+) = \text{ts}(\chi(k^+)) - \text{ts}(\chi(k))$. Applying this to all $i \leq k < j$ we get $\text{AC}(i, j) = \text{tt}$ iff $D(i, j) = \text{ts}(\chi(j)) - \text{ts}(\chi(i))$.

Finally, $\text{AC}(i, j) = \text{ff}$ implies $\text{ac}(k) = \text{ff}$ for some $i \leq k < j$. Using (A_4) we obtain $\text{ts}(\chi(j)) - \text{ts}(\chi(i)) \geq \text{ts}(\chi(k^+)) - \text{ts}(\chi(k)) \geq M$. \square

The transitions of $\mathcal{A}_{\text{valid}}^{K,M}$ are defined in Table 2.

Lemma 4.4. *Let τ be a (K, M) -STT and assume that $\mathcal{A}_{\text{valid}}^{K,M}$ has a run on τ reaching state q . Then, τ is monotonic and q is a realizable abstraction of τ .*

Proof. The conditions on the transitions for $\text{Rename}_{i,j}$, $\text{Add}_{i,j}^{\rightarrow}$ and $\text{Add}_{i,j}^{\curvearrowright I}$ directly ensure that the term is monotonic. We show that (A_1-A_4) are maintained by transitions of $\mathcal{A}_{\text{valid}}^{K,M}$.

- **Atomic STTs:** Consider a transition $\xrightarrow{(i,a)} q$ of $\mathcal{A}_{\text{valid}}^{K,M}$.
It is clear that q is a realizable abstraction of the atomic STT $\tau = (i, a)$.
- **Rename $_{i,j}$:** Consider a transition $q \xrightarrow{\text{Rename}_{i,j}} q'$ of $\mathcal{A}_{\text{valid}}^{K,M}$.
Assume that q is a realizable abstraction of some (K, M) -STT τ and let $\tau' = \text{Rename}_{i,j} \tau$. It is easy to check that q' is a realizable abstraction of τ' .
- **Forget $_i$:** Consider a transition $q \xrightarrow{\text{Forget}_i} q'$ of $\mathcal{A}_{\text{valid}}^{K,M}$.
Assume that q is a realizable abstraction of some (K, M) -STT τ and let $\tau' = \text{Forget}_i \tau$. It is easy to check that q' is a realizable abstraction of τ' . In particular, the correctness of the update $\text{ac}'(i^-)$ follows from Claim 4.3.
- **Add $_{i,j}^{\rightarrow}$:** Consider a transition $q \xrightarrow{\text{Add}_{i,j}^{\rightarrow}} q'$ of $\mathcal{A}_{\text{valid}}^{K,M}$.
Assume that q is a realizable abstraction of some (K, M) -STT τ and let $\tau' = \text{Add}_{i,j}^{\rightarrow} \tau$. It is easy to check that q' is a realizable abstraction of τ' .
- **Add $_{i,j}^{\curvearrowright I}$:** Consider a transition $q \xrightarrow{\text{Add}_{i,j}^{\curvearrowright I}} q'$ of $\mathcal{A}_{\text{valid}}^{K,M}$.
Assume that q is a realizable abstraction of some (K, M) -STT τ with the timestamp map ts . It is easy to check that q' is a realizable abstraction of $\tau' = \text{Add}_{i,j}^{\curvearrowright I} \tau$ with the same timestamp map $\text{ts}' = \text{ts}$. We only have to check the properties for the new timing constraint. First, condition $i < j$ and Claim 4.3 ensures that the new timing constraint is compatible with the linear order $<' = <$.
Second, by Claim 4.3, if $\text{AC}(i, j) = \text{ff}$ then the time elapsed between $\chi(i)$ and $\chi(j)$ is big and the timing constraint I is satisfied iff $I.\text{up} = \infty$. On the other hand, if $\text{AC}(i, j) = \text{tt}$ then using again Claim 4.3 we deduce that $D(i, j)$ is the actual time elapsed between $\chi(i)$ and $\chi(j)$ and the timing constraint I is satisfied iff $D(i, j) \in I$.
- **\oplus :** Consider a transition $q_1, q_2 \xrightarrow{\oplus} q$ of $\mathcal{A}_{\text{valid}}^{K,M}$.
Assume that q_1 and q_2 are realizable abstractions of some (K, M) -STTs τ_1 and τ_2 with timestamp maps ts_1 and ts_2 respectively. Let $\tau = \tau_1 \oplus \tau_2$. We show that q is a realizable abstraction of τ . Notice that $\llbracket \tau \rrbracket$ is the disjoint union of $\llbracket \tau_1 \rrbracket$ and $\llbracket \tau_2 \rrbracket$.
 (A_1) We have $\text{dom}(\chi) = \text{dom}(\chi_1) \cup \text{dom}(\chi_2) = P_1 \cup P_2 = P$.

(i, a)	$\xrightarrow{(i,a)} q = (P, \text{sd}, \text{tsm}, \text{ac})$ is a transition if $P = \{i\}$, $\text{sd}(i) = \text{ff}$ and $\text{ac}(i) = \text{ff}$. When reading an atomic STT $\tau = (i, a)$, only $\text{tsm}(i)$ is guessed.
$\text{Rename}_{i,j}$	$q = (P, \text{sd}, \text{tsm}, \text{ac}) \xrightarrow{\text{Rename}_{i,j}} q' = (P', \text{sd}', \text{tsm}', \text{ac}')$ is a transition if $i \in P$ and $i^- < j < i^+$. Then, q' is obtained from q by replacing i by j .
Forget_i	$q = (P, \text{sd}, \text{tsm}, \text{ac}) \xrightarrow{\text{Forget}_i} q' = (P', \text{sd}', \text{tsm}', \text{ac}')$ is a transition if $i^-, i, i^+ \in P$ and $\text{sd}(i^-) = \text{tt} = \text{sd}(i)$ (endpoints should stay colored). Then, state q' is deterministically given by $P' = P \setminus \{i\}$ and $\text{ac}'(i^-) = \text{AC}(i^-, i^+) \wedge (D(i^-, i^+) < M)$, the other values of $\text{sd}', \text{tsm}', \text{ac}'$ being inherited from $\text{sd}, \text{tsm}, \text{ac}$.
$\text{Add}_{i,j}^{\rightarrow}$	$q = (P, \text{sd}, \text{tsm}, \text{ac}) \xrightarrow{\text{Add}_{i,j}^{\rightarrow}} q' = (P, \text{sd}', \text{tsm}, \text{ac})$ is a transition if $i, j \in P$, $j = i^+$ and $\text{sd}(i) = \text{ff}$ ($\chi(i)$ does not already have a \rightarrow successor). The update is given by $\text{sd}'(i) = \text{tt}$ since we have added a \rightarrow -edge between $\chi(i)$ and $\chi(i^+)$, and all other values are unchanged.
$\text{Add}_{i,j}^{\wedge I}$	$q = (P, \text{sd}, \text{tsm}, \text{ac}) \xrightarrow{\text{Add}_{i,j}^{\wedge I}} q$ is a transition if $i, j \in P$, $i < j$ and either $(\text{AC}(i, j) = \text{tt}$ and $D(i, j) \in I)$ or $(\text{AC}(i, j) = \text{ff}$ and $I.\text{up} = \infty)$.
\oplus	<p>$q_1, q_2 \xrightarrow{\oplus} q$ where $q_1 = (P_1, \text{sd}_1, \text{tsm}_1, \text{ac}_1)$, $q_2 = (P_2, \text{sd}_2, \text{tsm}_2, \text{ac}_2)$ and $q = (P, \text{sd}, \text{tsm}, \text{ac})$ is a transition if the following hold</p> <ul style="list-style-type: none"> • $P_1 \cap P_2 = \emptyset$: the \oplus operation on STTs requires that the “active” colors of the two arguments are disjoint. • $P = P_1 \cup P_2$, $\text{sd} = \text{sd}_1 \cup \text{sd}_2$ and $\text{tsm} = \text{tsm}_1 \cup \text{tsm}_2$: these updates are deterministic. • the \rightarrow-blocks of the two arguments are shuffled according to the ordering of the colors. Formally, we check that it is not possible to insert a point from one argument inside a \rightarrow-block of the other argument: $\forall i \in P_1 \quad \text{sd}_1(i) \implies \text{next}_P(i) = \text{next}_{P_1}(i)$ $\forall i \in P_2 \quad \text{sd}_2(i) \implies \text{next}_P(i) = \text{next}_{P_2}(i).$ • Finally, ac satisfies $\text{ac}(\max(P)) = \text{ff}$ and $\forall i \in P_1 \setminus \{\max(P_1)\} \quad \text{ac}_1(i) \iff \text{AC}_q(i, \text{next}_{P_1}(i)) \wedge D_q(i, \text{next}_{P_1}(i)) < M$ $\forall i \in P_2 \setminus \{\max(P_2)\} \quad \text{ac}_2(i) \iff \text{AC}_q(i, \text{next}_{P_2}(i)) \wedge D_q(i, \text{next}_{P_2}(i)) < M.$
	<p>Notice that these conditions have several consequences.</p> <ul style="list-style-type: none"> • For all $i \in P_1$, if $\text{next}_P(i) = \text{next}_{P_1}(i)$ then $\text{ac}(i) = \text{ac}_1(i)$. For all $i \in P_2$, if $\text{next}_P(i) = \text{next}_{P_2}(i)$ then $\text{ac}(i) = \text{ac}_2(i)$. • For all $i \in P_1$, if $\text{next}_P(i) < \text{next}_{P_1}(i) = j < \infty$ (some points $\text{next}_P(i), \dots, \text{prev}_P(j)$ of the second argument have been inserted in the hole (i, j) of the first argument) then $\text{ac}_1(i) = \text{tt}$ implies $\text{ac}(i) = \text{ac}(\text{prev}_P(j)) = \text{tt}$ and $\text{AC}_{q_2}(\text{next}_P(i), \text{prev}_P(j)) = \text{tt}$. For all $i \in P_2$, if $\text{next}_P(i) < \text{next}_{P_2}(i) = j < \infty$ then $\text{ac}_2(i) = \text{tt}$ implies $\text{ac}(i) = \text{ac}(\text{prev}_P(j)) = \text{tt}$, $\text{AC}_{q_1}(\text{next}_P(i), \text{prev}_P(j)) = \text{tt}$.

Table 2: Transitions of $\mathcal{A}_{\text{valid}}^{K,M}$.

Moreover, $\text{EP}(\llbracket \tau \rrbracket) = \text{EP}(\llbracket \tau_1 \rrbracket) \cup \text{EP}(\llbracket \tau_2 \rrbracket) \subseteq \chi(P_1) \cup \chi(P_2) = \chi(P)$.

(A₂) Let $i \in P_1$ be such that $j = \text{next}_P(i) = \text{next}_{P_1}(i)$. We have $\text{sd}(i) = \text{tt}$ iff $\text{sd}_1(i) = \text{tt}$ iff $j \neq \infty$ and $\chi_1(i) \rightarrow^+ \chi_1(j)$ in $\llbracket \tau_1 \rrbracket$ iff $j \neq \infty$ and $\chi(i) \rightarrow^+ \chi(j)$ in $\llbracket \tau \rrbracket$.

Let $i \in P_1$ be such that $j = \text{next}_P(i) < \text{next}_{P_1}(i)$. We have $\text{sd}(i) = \text{sd}_1(i) = \text{ff}$, $j \neq \infty$ and there are no \rightarrow -path from $\chi(i) = \chi_1(i)$ to $\chi(j) = \chi_2(j)$ in $\llbracket \tau \rrbracket$ since $\llbracket \tau_1 \rrbracket$ and $\llbracket \tau_2 \rrbracket$ are disjoint.

We argue similarly for $i \in P_2$ which concludes the proof of Condition (A₂).

(A₃) Let $\dashrightarrow = \{(\chi(i), \chi(j)) \mid i \in P \wedge j = \text{next}_P(i) \neq \infty \wedge \text{sd}(i) = \text{ff}\}$. Let $< = (\rightarrow \cup \dashrightarrow)^+$. Using (A₂) and the definition of $<$, it is easy to see that for all $i, j \in P$, if $i < j$ then $\chi(i) < \chi(j)$. We deduce that $<_1 \cup <_2 \subseteq <$.

Let $u \curvearrowright^I v$ be a timing constraint in $\llbracket \tau \rrbracket$. Either it is in $\llbracket \tau_1 \rrbracket$ and it is compatible with $<_1$, hence also with $<$. Or it is in $\llbracket \tau_2 \rrbracket$ and it is compatible with $<_2$ and with $<$.

Next, we show that blocks of $\llbracket \tau_1 \rrbracket$ and $\llbracket \tau_2 \rrbracket$ do not overlap in $(\llbracket \tau \rrbracket, \dashrightarrow)$. Let i_1, j_1 be the colors of the left and right endpoints of some block $\chi(i_1) \rightarrow^* \chi(j_1)$ in $\llbracket \tau_1 \rrbracket$. For all $k \in P_1$ such that $i_1 \leq k < j_1$ we have $\chi(i_1) \rightarrow^* \chi(k) \rightarrow^+ \chi(\text{next}_{P_1}(k)) \rightarrow^* \chi(j_1)$. Applying (A₂) we get $\text{sd}_1(k) = \text{tt}$. Now, using the definition of the transition for \oplus , we get $\text{next}_P(k) = \text{next}_{P_1}(k)$. We deduce by induction that for all $\ell \in P_2$, if $i_1 < \ell$ then $j_1 < \ell$. By symmetry, the same holds for a block of $\llbracket \tau_2 \rrbracket$ and a color of P_1 .

Let i_2, j_2 be the colors of the left and right endpoints of some block $\chi(i_2) \rightarrow^* \chi(j_2)$ in $\llbracket \tau_2 \rrbracket$. Either $i_1 < i_2$ and we get $j_1 < i_2$ hence $\chi(j_1) < \chi(i_2)$. Or $i_2 < i_1$ and we get $j_2 < i_1$ hence $\chi(j_2) < \chi(i_1)$. We deduce that the blocks of $\llbracket \tau_1 \rrbracket$ and $\llbracket \tau_2 \rrbracket$ are shuffled in $(\llbracket \tau \rrbracket, \dashrightarrow)$ according to the order of the colors of their endpoints.

Therefore, $<$ is a total order on V and $(\llbracket \tau \rrbracket, \dashrightarrow)$ is a split-TCW.

(A₄) We construct the timestamp map ts for τ inductively on $V = V_1 \uplus V_2$ following the successor relation $\leq = \rightarrow \cup \dashrightarrow$. If $v = \min(V)$ is the first point of the split-TCW, we let

$$\text{ts}(v) = \begin{cases} \text{ts}_1(v) & \text{if } v \in V_1 \\ \text{ts}_2(v) & \text{if } v \in V_2. \end{cases}$$

Next, if $\text{ts}(u)$ is defined and $u \rightarrow v$ then we let

$$\text{ts}(v) = \begin{cases} \text{ts}(u) + \text{ts}_1(v) - \text{ts}_1(u) & \text{if } u \in V_1 \\ \text{ts}(u) + \text{ts}_2(v) - \text{ts}_2(u) & \text{if } u \in V_2. \end{cases}$$

Finally, if $\text{ts}(u)$ is defined and $u \dashrightarrow v$ then, with $i, j \in P$ being the colors of u and v ($\chi(i) = u$ and $\chi(j) = v$), we let

$$\text{ts}(v) = \begin{cases} \text{ts}(u) + d_q(i, j) & \text{if } \text{ac}(i) = \text{tt} \\ \text{ts}(u) + d_q(i, j) + M & \text{if } \text{ac}(i) = \text{ff}. \end{cases}$$

With this definition, the following hold

- Time is clearly non-decreasing: $\text{ts}(u) \leq \text{ts}(v)$ for all $u \leq v$
- (tsm, ac) is the modulo M abstraction of ts . The proof is by induction. First, if $i = \min(P)$ then $v = \min(V) = \chi(i)$ and $i \in P_1$ iff $v \in V_1$. Using the definitions of tsm and ts , we deduce easily that $\text{tsm}(i) = \text{ts}(v)[M]$. Next, let $i \in P$ with $j = \text{next}_P(i) < \infty$. Let $u = \chi(i)$, $v = \chi(j)$ and assume that $\text{tsm}(i) = \text{ts}(u)[M]$.

If $u \rightarrow^+ v$ and $u \in V_1$ then $v \in V_1$, $i, j \in P_1$, $\text{sd}_1(i) = \text{tt}$ and

$$\begin{aligned} \text{ts}(v)[M] &= (\text{ts}(u)[M] + \text{ts}_1(v)[M] - \text{ts}_1(u)[M])[M] \\ &= (\text{tsm}(i) + \text{tsm}_1(j) - \text{tsm}_1(i))[M] = \text{tsm}(j). \end{aligned}$$

Moreover, $\text{ac}(i) = \text{ac}_1(i)$ and $\text{ts}(v) - \text{ts}(u) = \text{ts}_1(v) - \text{ts}_1(u)$. We deduce that $\text{ac}(i) = \text{tt}$ iff $\text{ts}(v) - \text{ts}(u) < M$. The proof is similar if $u \rightarrow^+ v$ and $u \in V_2$.

Now, if $u \not\rightarrow^+ v$ then $u \dashrightarrow v$ (endpoints are always colored). We deduce that

$$\text{ts}(v)[M] = (\text{ts}(u)[M] + d_q(i, j))[M] = (\text{tsm}(i) + \text{tsm}(j) - \text{tsm}(i))[M] = \text{tsm}(j).$$

Moreover, it is clear that $\text{ts}(v) - \text{ts}(u) < M$ iff $\text{ac}(i) = \text{tt}$.

- Constraints are satisfied. Let $u \curvearrowright^I v$ be a timing constraint in $\llbracket \tau \rrbracket = \llbracket \tau_1 \rrbracket \uplus \llbracket \tau_2 \rrbracket$. Wlog we assume that $u, v \in V_1$. We know that $\text{ts}_1(v) - \text{ts}_1(u) \in I$.

If $u \rightarrow^+ v$ then we get $\text{ts}(v) - \text{ts}(u) = \text{ts}_1(v) - \text{ts}_1(u)$ from the definition of ts above. Hence, $\text{ts}(v) - \text{ts}(u) \in I$.

Now assume there are holes between u and v in $(\llbracket \tau \rrbracket, \dashrightarrow)$. Let u' be the right endpoint of the block of u and v' be the left endpoint of the block of v . Since endpoints are colored we find $i, j \in P$ such that $u' = \chi(i)$ and $v' = \chi(j)$. We have $i, j \in P_1$, $i < j$, $u \rightarrow^* u'$ and $v' \rightarrow^* v$. We deduce from the definition of ts that $\text{ts}(u') - \text{ts}(u) = \text{ts}_1(u') - \text{ts}_1(u)$ and $\text{ts}(v) - \text{ts}(v') = \text{ts}_1(v) - \text{ts}_1(v')$. Now, using Claim 4.5 below we obtain:

- * Either $\text{AC}_{q_1}(i, j) = \text{ff}$ and $\text{ts}(v') - \text{ts}(u') \geq M$. From Claim 4.3 we also have $\text{ts}_1(v') - \text{ts}_1(u') \geq M$. We deduce that $I \cdot \text{up} = \infty$ and $\text{ts}(v) - \text{ts}(u) \in I$.
- * Or $\text{AC}_{q_1}(i, j) = \text{tt}$ and $\text{ts}(v') - \text{ts}(u') = \text{ts}_1(v') - \text{ts}_1(u')$. Therefore,

$$\begin{aligned} \text{ts}(v) - \text{ts}(u) &= \text{ts}(v) - \text{ts}(v') + \text{ts}(v') - \text{ts}(u') + \text{ts}(u') - \text{ts}(u) \\ &= \text{ts}_1(v) - \text{ts}_1(v') + \text{ts}_1(v') - \text{ts}_1(u') + \text{ts}_1(u') - \text{ts}_1(u) \\ &= \text{ts}_1(v) - \text{ts}_1(u) \in I \end{aligned}$$

□

Claim 4.5. Let $i, j \in P_1$ with $i \leq j$ and let $u = \chi(i)$ and $v = \chi(j)$.

- (1) If $\text{AC}_{q_1}(i, j) = \text{ff}$ then $\text{ts}(v) - \text{ts}(u) \geq M$.
- (2) If $\text{AC}_{q_1}(i, j) = \text{tt}$ then $\text{ts}(v) - \text{ts}(u) = \text{ts}_1(v) - \text{ts}_1(u)$.

Proof. The proof is by induction on the number of points in P_1 between i and j . The result is clear if $i = j$. So assume that $k = \text{next}_{P_1}(i) \leq j$ and let $w = \chi(k)$. By induction, the claim holds for the pair (k, j) .

- (1) If $\text{AC}_{q_1}(i, j) = \text{ff}$ then either $\text{ac}_1(i) = \text{ff}$ or $\text{AC}_{q_1}(k, j) = \text{ff}$.
In the first case, by definition of the transition for \oplus , we have either $\text{AC}_q(i, k) = \text{ff}$ or $D_q(i, k) \geq M$. In both cases, we get $\text{ts}(w) - \text{ts}(u) \geq M$ by Claim 4.3.
In the second case, we get $\text{ts}(v) - \text{ts}(w) \geq M$ by induction.
Since ts is non-decreasing, we obtain $\text{ts}(v) - \text{ts}(u) \geq M$.
- (2) If $\text{AC}_{q_1}(i, j) = \text{tt}$ then $\text{ac}_1(i) = \text{tt}$ and $\text{AC}_{q_1}(k, j) = \text{tt}$.
By induction, we obtain $\text{ts}(v) - \text{ts}(w) = \text{ts}_1(v) - \text{ts}_1(w)$.
From the definition of the transition for \oplus , since $\text{ac}_1(i) = \text{tt}$, we get $\text{AC}_q(i, k) = \text{tt}$ and $D_q(i, k) < M$. Using Claim 4.3 we deduce that $\text{ts}(w) - \text{ts}(u) = D_q(i, k)$. Now, $D_q(i, k) < M$ implies $D_q(i, k) = d_q(i, k) = d_{q_1}(i, k)$. Using again Claim 4.3 we get $d_{q_1}(i, k) = \text{ts}_1(w) - \text{ts}_1(u)$. We conclude that $\text{ts}(w) - \text{ts}(u) = \text{ts}_1(w) - \text{ts}_1(u)$.

Combining the two equalities, we obtain $\text{ts}(v) - \text{ts}(u) = \text{ts}_1(v) - \text{ts}_1(u)$ as desired. \square

Accepting condition. The accepting states of $\mathcal{A}_{\text{valid}}^{K,M}$ should correspond to abstractions of TCWs. Hence the accepting states are of the form $(\{i\}, \text{sd}, \text{tsm}, \text{ac})$ (reached while reading an atomic term (i, a)) or $(\{i, j\}, \text{sd}, \text{tsm}, \text{ac})$ with $i, j \in \{1, \dots, K\}$, $i < j$, $\text{sd}(i) = \text{tt}$ and $\text{sd}(j) = \text{ff} = \text{ac}(j)$.

We show the correctness of the construction.

(\subseteq) Let τ be an STT accepted by $\mathcal{A}_{\text{valid}}^{K,M}$. There is an accepting run of $\mathcal{A}_{\text{valid}}^{K,M}$ reading τ and reaching state q at the root of τ . By Lemma 4.4, the term τ is monotonic and the state q is a realizable abstraction of τ , hence $(\llbracket \tau \rrbracket, \dashrightarrow)$ is a split-TCW. But since q is accepting, we have $\dashrightarrow = \emptyset$. Hence $\llbracket \tau \rrbracket$ is a TCW. Moreover, from (A₄) we deduce that $\llbracket \tau \rrbracket$ is realizable and the endpoints of $\llbracket \tau \rrbracket$ are the only colored points by (A₁) and the acceptance condition.

(\supseteq) Let τ be a monotonic (K, M) -STTs such that $\llbracket \tau \rrbracket = (G, \chi)$ is a realizable TCW and the endpoints of $\llbracket \tau \rrbracket$ are the only colored points. Let $\text{ts}: V \rightarrow \mathbb{N}$ be a timestamp map satisfying all the timing constraints in τ . We construct a run of $\mathcal{A}_{\text{valid}}^{K,M}$ on τ by resolving the non-deterministic choices as explained below. Notice that the transitions for $\text{Rename}_{i,j}$, Forget_i , $\text{Add}_{i,j}^{\rightarrow}$ and $\text{Add}_{i,j}^{\wedge^I}$ are deterministic. We will obtain an accepting run ρ of $\mathcal{A}_{\text{valid}}^{K,M}$ on τ such that for every subterm τ' , the state $\rho(\tau')$ satisfies (A₄) with timestamp map ts , or more precisely, with the restriction of ts to the vertices in $\llbracket \tau' \rrbracket$.

- A leaf (i, a) of the term τ corresponds to some vertex $v \in V$. The transition taken at this leaf will set $\text{tsm}(i) = \text{ts}(v)[M]$ so that (A₄) holds with ts .
- We can check that the conditions enabling transitions at $\text{Rename}_{i,j}$, Forget_i or $\text{Add}_{i,j}^{\rightarrow}$ nodes are satisfied since τ is monotonic and $\llbracket \tau \rrbracket$ is a TCW whose endpoints are colored.
- For a subterm $\tau' = \text{Add}_{i,j}^{\wedge^I} \tau''$ the condition $i < j$ is satisfied since τ is monotonic and the condition involving AC and D is satisfied by Claim 4.3 since ts satisfies all timing constraints of τ and (A₄) holds with ts at τ'' .
- Consider a subterm $\tau' = \tau_1 \oplus \tau_2$. Let $\rho(\tau_1) = q_1 = (P_1, \text{sd}_1, \text{tsm}_1, \text{ac}_1)$ and $\rho(\tau_2) = q_2 = (P_2, \text{sd}_2, \text{tsm}_2, \text{ac}_2)$. The active colors of τ_1 and τ_2 are disjoint, hence $P_1 \cap P_2 = \emptyset$ by (A₁). Define $q' = (P', \text{sd}', \text{tsm}', \text{ac}')$ by $P' = P_1 \cup P_2$, $\text{sd}' = \text{sd}_1 \cup \text{sd}_2$, $\text{tsm}' = \text{tsm}_1 \cup \text{tsm}_2$ and for all $i \in P'$, $\text{ac}'(i) = \text{tt}$ iff $i^+ \neq \infty$ and $\text{ts}(\chi'(i^+)) - \text{ts}(\chi'(i)) < M$.

The condition for $q_1, q_2 \xrightarrow{\oplus} q'$ on shuffling the \rightarrow -blocks holds since $\llbracket \tau \rrbracket$ is a TCW. Now, we look at the condition on ac' . Let $i \in P_1 \setminus \{\max(P_1)\}$ and $j = \text{next}_{P_1}(i)$. We have $\text{ac}_1(i) = \text{tt}$ iff $\text{ts}(\chi_1(j)) - \text{ts}(\chi_1(i)) < M$ since (A₄) holds with ts at τ_1 . The latter holds iff for all $k \in P'$ with $i \leq k < j$ we have $\text{ts}(\chi'(k^+)) - \text{ts}(\chi'(k)) < M$ (i.e., $\text{AC}_{q'}(i, j) = \text{tt}$ by the above definition of ac') and $D_{q'}(i, j) < M$ (again, by the definition of ac' we have that $\text{ac}'(k) = \text{tt}$ implies $d_{q'}(k, k^+) = \text{ts}(\chi'(k^+)) - \text{ts}(\chi'(k))$ and $\text{AC}'(i, j) = \text{tt}$ implies $D_{q'}(i, j) = \text{ts}(\chi'(j)) - \text{ts}(\chi'(i))$). \square

To ensure that we accept only terms denoting simple TCWs of split-width at most K , we adapt our above construction to $\mathcal{A}_{\text{valid}}^{2K,M}$. We will refer this restriction to $\mathcal{A}_{\text{valid}}^{2K,M}$ as $\mathcal{B}_{\text{valid}}^{K,M}$.

Corollary 4.6. *We can build a tree automaton $\mathcal{B}_{\text{valid}}^{K,M}$ with $M^{\mathcal{O}(K)}$ states such that $\tau \in \mathcal{L}(\mathcal{B}_{\text{valid}}^{K,M})$ iff $\llbracket \tau \rrbracket \in \text{STCW}^{K,M}$ is realizable.*

Proof. We will construct $\mathcal{B}_{\text{valid}}^{K,M}$ as a restriction of $\mathcal{A} = \mathcal{A}_{\text{valid}}^{2K,M}$.

Let τ be a term accepted by \mathcal{A} . Notice that, for every subterm τ' of τ , the number of blocks in $\llbracket \tau' \rrbracket$ can be computed from the state $q' = (P', \text{sd}', \text{tsm}', \text{ac}')$ that labels τ' in this accepting run: this number of blocks is $|\{i \in P' \mid \text{sd}(i) = \text{ff}\}|$. So we restrict \mathcal{A} to the set of states q' such that the number of blocks of q' is at most K . The automaton $\mathcal{B}_{\text{valid}}^{K,M}$ is defined by further restricting \mathcal{A} so that adding a timing constraint $\text{Add}_{i,j}^I$ is allowed only on subterms of the form $(i, a) \oplus (j, b)$ with $i < j$.

Then, an accepted term $\tau \in \mathcal{L}(\mathcal{B}_{\text{valid}}^{K,M})$ is realizable since it is accepted by \mathcal{A} (since $\mathcal{B}_{\text{valid}}^{K,M}$ is a restriction of \mathcal{A}) and it describes a split-decomposition of $\llbracket \tau \rrbracket = (\mathcal{V}, \chi)$ of width at most K . Further, since timing constraints are added only on terms of the form $(i, a) \oplus (j, b)$, the TCW $\llbracket \tau \rrbracket$ must be simple. We obtain $\llbracket \tau \rrbracket \in \text{STCW}^{K,M}$ and is realizable.

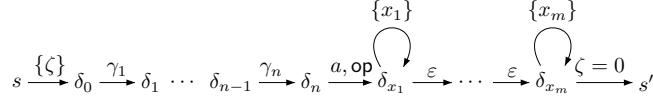
Conversely, let \mathcal{V} be a realizable simple TCW in $\text{STCW}^{K,M}$. There is a split-tree T of width at most K for \mathcal{V} . By Lemma 3.1, there is a $(2K, M)$ -STT τ with $\llbracket \tau \rrbracket = (\mathcal{V}, \chi)$. We may assume that τ is monotonic and only the endpoints of \mathcal{V} are colored by χ . By Theorem 4.2, we get $\tau \in \mathcal{L}(\mathcal{A})$. Moreover, for every subterm τ' of τ , the graph $\llbracket \tau' \rrbracket$ has at most K blocks. Also, the timing constraints in τ are only added by subterms of the form $\text{Add}_{i,j}^I((i, a) \oplus (j, b))$ with $i < j$. Therefore, $\tau \in \mathcal{L}(\mathcal{B}_{\text{valid}}^{K,M})$. \square

5. TREE AUTOMATA FOR TIMED SYSTEMS

The goal of this section is to build a tree automaton which accepts the STTs denoting (realizable) STCWs accepted by a TPDA. Again, to prove the existence of a tree automaton, it suffices to show the MSO-definability of the existence of a (realizable) run of \mathcal{S} on a simple TCW, and appeal to Courcelle's theorem [9]. However, as explained in the previous section, this may not give an optimal complexity. Hence, we choose to directly construct the tree automaton $\mathcal{A}_{\mathcal{S}}^{K,M}$, where M is one more than the maximal constant used in the TPDA, and K is the split width. Our input is therefore the system, in this case a TPDA \mathcal{S} , whose size $|\mathcal{S}|$ includes the number of states, transitions and the maximal constant used in \mathcal{S} .

Let us explain first how the automaton would work for an *untimed* system with no stack. At a leaf of the STT of the form (i, a) , the tree automaton guesses a transition δ from \mathcal{S} which could be executed reading action a . It keeps the transition in its state, paired with color i . After reading a subterm τ , the tree automaton stores in its state a map Δ which assigns with each active (free) color k of τ the transition $\Delta(k)$ guessed at the corresponding leaf. The map Δ can be easily updated at nodes labelled \oplus or **Forget** _{i} or **Rename** _{i,j} . When, reading a node labelled $\text{Add}_{i,j}^I$, the tree automaton checks that the target state of the transition $\Delta(i)$ equals the source state of the transition $\Delta(j)$. This ensures that the transitions guessed at leaves form a run when taken in the total order induced by the word denoted by the final term. At the root, we check that at most two colors remain free: i and j for the leftmost (resp. rightmost) endpoint of the word. Then, the tree automaton accepts if the source state of $\Delta(i)$ is initial in \mathcal{S} and the target state of $\Delta(j)$ is final in \mathcal{S} .

The situation is a bit more complicated for timed systems. First, we are interested in the *simple* TCW semantics in which each event is blown-up in several micro-events. Following this idea, each transition $(s, \gamma, a, \text{op}, R, s')$ of the timed system \mathcal{S} is blown-up in micro-transitions as explained in Section 2.3, assuming that $\gamma = \gamma_1 \wedge \dots \wedge \gamma_n$ has n conjuncts of the form $x \in I$, and that $R = \{x_1, x_2, \dots, x_m\}$:



Notice the reset part at the end which allows an arbitrary number (possibly zero) resets of clock x_1 (depending on how many timing constraints for clock x_1 will originate from this reset), followed by an arbitrary number (possibly zero) resets of clock x_2 , etc.

The second difficulty is to make sure that, when a guard of the form $x \in I$ is checked in some transition, then the source point of the timing constraint in the simple TCW indeed corresponds to the latest transition resetting clock x . To check this property, the tree automaton $\mathcal{A}_S^{K,M}$ stores,

- for each color k of the left endpoint of a block, the set $C(k)$ of clocks reset by some transition whose *middle* point (the one corresponding to a, op micro-transition) occurs in the block.
- for each pair (i, j) of colors of left endpoints of *distinct* blocks, the set $D(i, j)$ of clocks that are reset in block i and checked in j .

Finally, the last property to be checked is that the push-pop edges are well-nested. To this end, the tree automaton $\mathcal{A}_S^{K,M}$ stores the set E of pairs (i, j) of left endpoints of *distinct* blocks such that there is at least one push-pop edge from block i to block j .

Formally, a state of $\mathcal{A}_S^{K,M}$ is a tuple (q, Δ, C, D, E) where $q = (P, \dots)$ is a state of $\mathcal{B}_{\text{valid}}^{K,M}$ (as seen in Theorem 4.2) the map Δ assigns to each color $k \in P$ the pair of states of the micro-transition $\Delta(k)$ guessed at the leaf corresponding to color k , and the maps C, D and the set E are as described above. The number of states of $\mathcal{A}_S^{K,M}$ is hence calculated as follows:

- q depends on P ; the number of possible maps P is $2^{\mathcal{O}(K)}$.
- The size of Δ is $[S^2]^K$, where S denotes the states of the underlying system \mathcal{S} .
- The size of C is $[2^{|X|}]^K$, while the size of D is $[2^{|X|}]^{K^2}$, and the size of E is 2^{K^2} .

The number of states of $\mathcal{A}_S^{K,M}$ is hence bounded by $|\mathcal{S}|^{\mathcal{O}(K)} \cdot 2^{\mathcal{O}(K^2(|X|+1))}$.

The transitions of $\mathcal{A}_S^{K,M}$ will check the existence of transitions of $\mathcal{B}_{\text{valid}}^{K,M}$ to check for validity and in addition will check the existence of an abstract run of the system \mathcal{S} . These are described in detail in Tables 3, 4. With this, we finally have the acceptance condition as follows.

A state (q, Δ, C, D, E) is accepting if

- q is an accepting state of $\mathcal{B}_{\text{valid}}^{K,M}$, hence it consists of a single block with left endpoint i and right endpoint j (possibly $i = j$),
- the source state of $\Delta(i)$ is an initial state of \mathcal{S} , and the target state of $\Delta(j)$ is a final state of \mathcal{S} , or there is an ε -path of micro-transitions from the target state of $\Delta(j)$ to some final state of \mathcal{S} .

If the underlying system is a timed automaton, it is sufficient to store the maps C, D in the state, since there are no stack operations. From the above construction we obtain the following theorem:

Theorem 5.1. *Let \mathcal{S} be a TPDA of size $|\mathcal{S}|$ with set of clocks X and using constants less than M . Then, we can build a tree automaton $\mathcal{A}_S^{K,M}$ of size $|\mathcal{S}|^{\mathcal{O}(K)} \cdot 2^{\mathcal{O}(K^2(|X|+1))}$ such that $\mathcal{L}(\mathcal{A}_S^{K,M}) = \{\tau \in \mathcal{L}(\mathcal{B}_{\text{valid}}^{K,M}) \mid \llbracket \tau \rrbracket \in \text{STCW}(\mathcal{S})\}$.*

(i, a)	$\xrightarrow{(i,a)} (q, \Delta, C, D, E)$ is a transition of $\mathcal{A}_S^{K,M}$ if $\xrightarrow{(i,a)} q$ is a transition of $\mathcal{B}_{\text{valid}}^{K,M}$ and <ul style="list-style-type: none"> • either there is a transition $\delta = (s, \gamma, a, \text{nop}, R, s')$ of \mathcal{S} (the tree automaton $\mathcal{A}_S^{K,M}$ guesses this transition) such that: <ul style="list-style-type: none"> – $\Delta(i) = (\delta_n, \delta_x)$ is the <i>middle</i> micro-transition of δ (assuming γ has n conjuncts), – $C(i) = R$ since the set of clocks reset by δ is R, and – D is nowhere defined and $E = \emptyset$ since we have a single block. • or $a = \varepsilon$ and we take a special initial dummy transition $\delta = (s_{\text{dummy}}, \text{tt}, \varepsilon, \text{nop}, X, s_0)$, where s_0 is an initial state of \mathcal{S} and all clocks are reset.
$\text{Rename}_{i,j}$	$(q, \Delta, C, D, E) \xrightarrow{\text{Rename}_{i,j}} (q', \Delta', C', D', E')$ is a transition of $\mathcal{A}_S^{K,M}$ if $q \xrightarrow{\text{Rename}_{i,j}} q'$ is a transition of $\mathcal{B}_{\text{valid}}^{K,M}$ and Δ', C', D', E' are obtained from Δ, C, D, E , respectively, by replacing i by j .
Forget_i	$(q, \Delta, C, D, E) \xrightarrow{\text{Forget}_i} (q', \Delta', C', D', E')$ is a transition of $\mathcal{A}_S^{K,M}$ if $q \xrightarrow{\text{Forget}_i} q'$ is a transition of $\mathcal{B}_{\text{valid}}^{K,M}$ (in particular, i is not an endpoint) and Δ' is obtained from Δ by forgetting the entry of color i . Since i is not a left endpoint, C, D and E are not affected: $C' = C, D' = D$ and $E' = E$.
$\text{Add}_{i,j}^I((i, \varepsilon) \oplus (j, \varepsilon))$	$\xrightarrow{\text{Add}_{i,j}^I((i, \varepsilon) \oplus (j, \varepsilon))} (q, \Delta, C, D, E)$ is a transition of $\mathcal{A}_S^{K,M}$ if $\xrightarrow{\text{Add}_{i,j}^I((i, \varepsilon) \oplus (j, \varepsilon))} q$ is a transition of $\mathcal{B}_{\text{valid}}^{K,M}$ and either <ul style="list-style-type: none"> • there exists a clock constraint $x \in I$ induced by two transitions $\delta^1 = (s^1, \gamma^1, a^1, \text{op}^1, R^1, s'^1)$ and $\delta^2 = (s^2, \gamma^2, a^2, \text{op}^2, R^2, s'^2)$ of \mathcal{S}, i.e., $x \in R^1$ and some conjunct of γ^2, say the k-th, is $x \in I$ (the tree automaton $\mathcal{A}_S^{K,M}$ guesses this) such that: <ul style="list-style-type: none"> – $\Delta(i) = (\delta_x^1, \delta_x^1)$ is the reset micro-transition for clock x of δ^1, $\Delta(j) = (\delta_{k-1}^2, \delta_k^2)$ is the micro-transition checking the k-th conjunct of γ^2, – $C(i) = C(j) = \emptyset, D(i, j) = \{x\}$ and $E = \emptyset$. • or $I = [0, 0]$ and the tree automaton $\mathcal{A}_S^{K,M}$ guesses that it encodes the ζ clock constraint of some transition $\delta = (s, \gamma, a, \text{op}, R, s')$ of \mathcal{S}. Then, $\Delta(i) = (s, \delta_0), \Delta(j) = (\delta_x, s')$ are the first and last micro-transitions from δ, $C(i) = C(j) = \emptyset, D(i, j) = \emptyset$ and $E = \emptyset$.

Table 3: Transitions of $\mathcal{A}_S^{K,M}$ – part 1

Proof. Let τ be a K -STT and $\mathcal{V} = \llbracket \tau \rrbracket$. We will show that τ is accepted by $\mathcal{A}_S^{K,M}$ iff $\mathcal{V} \in \mathcal{L}(\mathcal{B}_{\text{valid}}^{K,M})$ and $\mathcal{V} \in \text{STCW}(\mathcal{S})$.

$\text{Add}_{i,j}^I((i, a^1) \oplus (j, a^2))$	$\frac{\text{Add}_{i,j}^I((i, a^1) \oplus (j, a^2))}{\text{Add}_{i,j}^I((i, a^1) \oplus (j, a^2))} \rightarrow (q, \Delta, C, D, E) \text{ is a transition of } \mathcal{A}_S^{K,M} \text{ if}$ $\frac{\text{Add}_{i,j}^I((i, a^1) \oplus (j, a^2))}{\text{Add}_{i,j}^I((i, a^1) \oplus (j, a^2))} \rightarrow q \text{ is a transition of } \mathcal{B}_{\text{valid}}^{K,M} \text{ and there are two match-}$ <p>ing push-pop transitions of \mathcal{S}: $\delta^1 = (s^1, \gamma^1, a^1, \downarrow_b, R^1, s'^1)$ and $\delta^2 = (s^2, \gamma^2, a^2, \uparrow_b^I, R^2, s'^2)$ (these transitions are guessed by $\mathcal{A}_S^{K,M}$) such that $\Delta(i) = (\delta_n^1, \delta_x^1)$ is the <i>middle</i> micro-transition of δ^1, $\Delta(j) = (\delta_m^2, \delta_y^2)$ is the <i>middle</i> micro-transition of δ^2, $C(i) = R^1$, $C(j) = R^2$, $D(i, j) = \emptyset$ and $E = \{(i, j)\}$.</p>
$\text{Add}_{i,j}^{\rightarrow}$	$(q, \Delta, C, D, E) \xrightarrow{\text{Add}_{i,j}^{\rightarrow}} (q', \Delta', C', D', E') \text{ is a transition of } \mathcal{A}_S^{K,M} \text{ if}$ $q \xrightarrow{\text{Add}_{i,j}^{\rightarrow}} q' \text{ is a transition of } \mathcal{B}_{\text{valid}}^{K,M} \text{ (thus, } i \text{ is a right endpoint and } j = i^+ \text{ is a left endpoint) and}$ <ul style="list-style-type: none"> • Either the target state of $\Delta(i)$ equals the source state of $\Delta(j)$, or there is an ε-path of micro-transitions $\xrightarrow{\varepsilon} \delta_{x_k} \xrightarrow{\varepsilon} \dots \xrightarrow{\varepsilon} \delta_{x_\ell} \xrightarrow{\varepsilon}$ between the target state of $\Delta(i)$ and the source state of $\Delta(j)$. Indeed, adding \rightarrow between i and j means that these points are consecutive in the final TCW, hence it should be possible to concatenate the micro-transition taken at i and j. • $\Delta' = \Delta$. • Let k be the left end-point of $\text{Block}(i)$. We merge $\text{Block}(k)$ and $\text{Block}(j)$ hence we set $C'(k) = C(k) \cup C(j)$ and $C'(\ell) = C(\ell)$ if $\ell \notin \{k, j\}$ is another left endpoint. Also, if $\ell, \ell' \notin \{k, j\}$ are other left endpoints, we set $D'(\ell, \ell') = D(\ell, \ell')$ if $\ell < \ell'$, $D'(\ell, k) = D(\ell, k) \cup D(\ell, j)$ if $\ell < k$, and $D'(k, \ell) = D(k, \ell) \cup D(j, \ell)$ if $j < \ell$. Finally, E' is the set of pairs (ℓ, ℓ') of left endpoints of distinct blocks in q' s.t. either $(\ell, \ell') \in E$, or $\ell = k$ and $(j, \ell') \in E$, or $\ell < k = \ell'$ and $(\ell, j) \in E$.
\oplus	$(q_1, \Delta_1, C_1, D_1, E_1), (q_2, \Delta_2, C_2, D_2, E_2) \xrightarrow{\oplus} (q, \Delta, C, D, E) \text{ is a transition of } \mathcal{A}_S^{K,M} \text{ if } q_1, q_2 \xrightarrow{\oplus} q \text{ is a transition of } \mathcal{B}_{\text{valid}}^{K,M} \text{ and}$ <ul style="list-style-type: none"> • The other components are inherited: $\Delta = \Delta_1 \cup \Delta_2$, $C = C_1 \cup C_2$, $D = D_1 \cup D_2$ and $E = E_1 \cup E_2$. • For all $i < k < j$ left endpoints in q, we have $C(k) \cap D(i, j) = \emptyset$. This is the crucial condition which ensures that a timing constraint always refers to the last transition resetting the clock being checked. If some clock $x \in D(i, j)$ is reset in $\text{Block}(i)$ and checked in $\text{Block}(j)$, it is not possible to insert $\text{Block}(k)$ between blocks of i and j if clock x is reset by a transition in $\text{Block}(k)$. • For all $(i, j) \in E$ and $(k, \ell) \in E$, if $i < k < j$ then $\ell \leq j$. This ensures that the push-pop edges are well-nested.

Table 4: Transitions of $\mathcal{A}_S^{K,M}$ – part 2

Assume that $\mathcal{A}_S^{K,M}$ has an accepting run on τ . Each move of $\mathcal{A}_S^{K,M}$ first checks if there is a move on $\mathcal{B}_{\text{valid}}^{K,M}$, hence we obtain an accepting run of $\mathcal{B}_{\text{valid}}^{K,M}$ on τ . Thus, $\mathcal{V} \in \mathcal{L}(\mathcal{B}_{\text{valid}}^{K,M})$ and so by Corollary 4.6 of Theorem 4.2, \mathcal{V} is a realizable STCW. It remains to check that \mathcal{V} is generated or accepted by \mathcal{S} . That is, we need to show that there exists an abstract accepting run of \mathcal{S} on \mathcal{V} which starts with an initial state, ends in a final state and satisfies the three conditions stated in Section 2.3.

- (1) We first define the sequence of transitions. Each vertex v of \mathcal{V} labeled by letter a is introduced as a node colored i in some atomic term (i, a) . We let $\delta(v) = \delta(i)$ be the transition guessed by $\mathcal{A}_S^{K,M}$ when reading this atomic term. We start by reading ε and guessing the special initial transition that resets all clocks and goes to an initial state s_0 . Subsequently, for each $u \rightarrow v$ in \mathcal{V} we will have an $\text{Add}_{i,j}^{\rightarrow}$ occurring in τ such that $\delta(u) = \delta(i)$ and $\delta(v) = \delta(j)$. By definition of $\text{Add}_{i,j}^{\rightarrow}$ transition, we will check at this point that either $\text{target}(\delta(u)) = \text{source}(\delta(v))$ or there is a sequence of ε micro-transitions between these two states : that is, we will check if the target of $\Delta(i)$ is the same as the source of $\Delta(j)$. So we have constructed a sequence of transitions $(\delta(v))_v$ which forms a path in \mathcal{S} reading \mathcal{V} and starting from the initial state s_0 . By the acceptance condition, if v is the maximal vertex of \mathcal{V} then $\text{target}(\delta(v))$ is a final state.
- (2) Now, we need to check that the sequence of push-pop operations is well-nested. This is achieved by using the set E which, as mentioned earlier, stores the pairs (i, j) of left endpoints of distinct blocks such that there is at least one push-pop edge from block i to block j . Note that every such push-pop edge gets into the set E by the last atomic rule in Table 3, where it is checked that the transitions that were guessed at i and j were indeed, matching push-pop transitions. This information is propagated correctly when we add a process edge by the last condition in $\text{Add}_{i,j}^{\rightarrow}$ rule. That is, if the adding of an edge changes the left end-point, then the elements of E are updated to refer to the new left end-point. This ensures that the information is not forgotten since (left) end-points are never forgotten. Now, observe that the only operation that may allow the push-pop edges to cross is the general combine. But the last condition, in the definition of transition here, uses the set E information to make sure that well-nesting is not violated. More precisely, if $(i, j), (k, \ell)$ is in $E_1 \cup E_2$ and $i < k < j$, we check that $\ell \leq j$ and hence the edges are well-nested.
- (3) Finally, we check that the clock constraints are properly matched (i.e., when a guard is checked, the source of the timing constraint is indeed the latest transition resetting this clock. To do this, we make use of the sets C and D . The set C maintains for each left-endpoint of a block, the set of clocks reset by transitions whose middle-points occur at that block. This is introduced at the atomic transition rule and maintained during all transitions and propagated correctly when left-endpoints change, which happens only when blocks are merged by $\text{Add}_{i,j}^{\rightarrow}$ rule (see last condition in this rule). The set $D(i, j)$ refers to the set of clocks that are reset in $\text{Block}(i)$ and checked at $\text{Block}(j)$. Again, this set is populated at an atomic matching rule of the form $\text{Add}_{i,j}^{\wedge I}((i, \varepsilon) \oplus (j, \varepsilon))$ as defined. And as before, it is maintained throughout and propagated when end-points of blocks change, i.e., during the adding of a process edge.

Now, for any guard $x \in I$ that is checked, say matching vertex u to v , we claim that (i) x is reset at u and checked at v wrt $x \in I$ and (ii) between u and v , x is

never reset. (i) follows from the definition of the atomic rule for adding a matching edge relation. And for (ii), we observe that once the matching relation is added between u, v , say with colors i, j resp., then $x \in D(i, j)$. If the general combine is not used, then we cannot add a reset in between, and hence, this holds. When the general combine is used, then we check whether for all $i < k < j$ with left end-points in q_1 or q_2 , whether $C(k) \cap (D(i, j)) = \emptyset$. This ensures that any block k inserted between i and j ($x \in D(i, j)$) means that there is a matching edge between blocks of i, j cannot contain a transition which resets the clock x (x is stored in $C(k)$). Thus, the last reset-point is preserved and our checks are indeed correct.

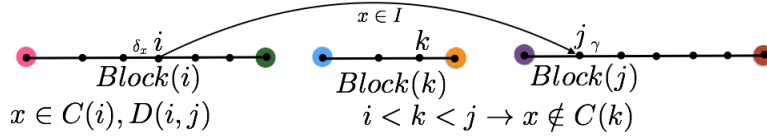


Figure 9: If x is reset in $\text{Block}(i)$, and checked in $\text{Block}(j)$, then $x \in D(i, j)$, and $x \in C(i)$. If $i < k < j$ is a point, then x cannot be reset in $\text{Block}(k)$, and $\text{Block}(k)$ cannot be inserted between $\text{Block}(i)$ and $\text{Block}(j)$.

Thus, we obtain that \mathcal{V} is indeed generated by \mathcal{S} , i.e., $\mathcal{V} \in \text{STCW}(\mathcal{S})$.

In the reverse direction, if $\mathcal{V} \in \text{STCW}(\mathcal{S}) \cap \mathcal{L}(\mathcal{B}_{\text{valid}}^{K,M})$, then there is a sequence of transitions which lead to the accepting state on reading \mathcal{V} in \mathcal{S} . By guessing each of these transitions correctly at every point (at the level of the atomic transitions), we can easily generate the run of our automaton $\mathcal{A}_{\mathcal{S}}^{K,M}$ which is also consequently accepting. \square

As a corollary we obtain,

Corollary 5.2. *Given a timed (pushdown) automaton \mathcal{S} , $\mathcal{L}(\mathcal{S}) \neq \emptyset$ iff $\mathcal{L}(\mathcal{A}_{\mathcal{S}}^{K,M}) \neq \emptyset$.*

Proof. If $\mathcal{L}(\mathcal{S}) \neq \emptyset$, then there is a realizable STCW \mathcal{V} accepted by \mathcal{S} . By Theorem 3.4, we know that its split-width is bounded by a constant $K = 4|X| + 6$. Further, from the proof of Theorem 3.4, we obtain that Eve has a winning strategy on \mathcal{V} with atmost $K - 1$ holes. Eve's strategy is such that it generates a K -STT τ such that $\mathcal{V} = \llbracket \tau \rrbracket$. Since $\mathcal{V} = \llbracket \tau \rrbracket$ is realizable, by Theorem 4.2, $\tau \in \mathcal{L}(\mathcal{B}_{\text{valid}}^{K,M})$. Thus, \mathcal{V} in $\text{STCW}(\mathcal{S})$ and $\tau \in \mathcal{L}(\mathcal{B}_{\text{valid}}^{K,M})$. Theorem 5.1 above then gives $\mathcal{V} \in \mathcal{L}(\mathcal{A}_{\mathcal{S}}^{K,M})$. The converse argument is similar. \square

As a consequence of all the above results, we obtain,

Theorem 5.3. *Checking emptiness for TPDA \mathcal{S} is decidable in EXPTIME, while timed automata are decidable in PSPACE.*

Proof. We first discuss the case of TPDA. First note that all STCWs in the semantics of a TPDA have a split-width bounded by some constant K . By Corollary 5.2, checking non-emptiness of \mathcal{S} boils down to checking non-emptiness of tree automata $\mathcal{A}_{\mathcal{S}}^{K,M}$. But while checking non-emptiness of $\mathcal{A}_{\mathcal{S}}^{K,M}$, the construction also appeals to $\mathcal{B}_{\text{valid}}^{K,M}$ and checks that its non-emptiness too. From the fact that checking emptiness of tree automata is in PTIME and given the sizes of the constructed tree automata $\mathcal{B}_{\text{valid}}^{K,M}$, $\mathcal{A}_{\mathcal{S}}^{K,M}$ (both exponential in the size of the input TPDA \mathcal{S}), we obtain an EXPTIME procedure to check emptiness of \mathcal{S} .

If the underlying system \mathcal{S} was a timed automaton, then as seen in Figure 8, Eve's strategy gives rise to a word-like decomposition of the STCW. Hence, the split-trees are

actually word-like binary trees, where one child is always an atomic node. Thus, we can use word automata (to guess the atomic node) instead of tree automata to check emptiness. Using the NLOGSPACE complexity of emptiness for word automata, and using the exponential sizes of $\mathcal{B}_{\text{valid}}^{K,M}$, $\mathcal{A}_{\mathcal{S}}^{K,M}$, we obtain a PSPACE procedure to check emptiness of \mathcal{S} . \square

6. DENSE TIME MULTI-STACK PUSHDOWN SYSTEMS

As another application of our technique, we now consider the model of dense-timed multi-stack pushdown automata (dtMPDA), which have several stacks. The reachability problem for untimed multi-stack pushdown automata (MPDA) is already undecidable, but several restrictions have been studied on (untimed) MPDA, like bounded rounds [14], bounded phase, bounded scope and so on to regain decidability.

In this section, we consider dtMPDA with the restriction of “bounded rounds”. To the best of our knowledge, this timed model has not been investigated until now. Our goal is to illustrate how our technique can easily be applied here with a minimal overhead (in difficulty and complexity).

Formally, a dtMPDA is a tuple $\mathcal{S} = (S, \Sigma, \Gamma, X, s_0, F, \Delta)$ similar to a TPDA defined in Section 2.3. The only difference is the stack operation op which now specifies which stack is being operated on. That is,

- (1) nop does not change the contents of any stack (same as before),
- (2) \downarrow_c^i where $c \in \Gamma$ is a push operation that adds c on top of stack i , with age 0.
- (3) $\uparrow_{c \in I}^i$ where $c \in \Gamma$ and $I \in \mathcal{I}$ is a pop operation that removes the top most symbol of stack i provided it is a c with age in the interval I .

A sequence $\sigma = \text{op}_1 \cdots \text{op}_m$ of operations is a *round* if it can be decomposed in $\sigma = \sigma_1 \cdots \sigma_n$ where each factor σ_i is a possibly empty sequence of operations of the form nop , \downarrow_c^i , $\uparrow_{c \in I}^i$.

Let us fix an integer bound k on the number of rounds. The semantics of the dtMPDA in terms of STCWs is exactly the same as for TPDA, except that the sequence of stack operations along any run is restricted to (at most) k rounds. Thus, any run of dtMPDA can be broken into a finite number of contexts, such that in each context only a single stack is used. As before, the sequence of push-pop operations of any stack must be well-nested.

Simple TC-word semantics for dtMPDA. We define the semantics for dtMPDA in terms of simple TCWs. Let n denote the number of stacks. A simple TCW $\mathcal{V} = (V, \rightarrow, (\curvearrowright^I)_{I \in \mathcal{I}}, \lambda)$ is said to be generated or accepted by a dtMPDA \mathcal{S} if there is an accepting abstract run $\rho = (s_0, \gamma_1, a_1, \text{op}_1, R_1, s_1) (s_1, \gamma_2, a_2, \text{op}_2, R_2, s_2) \cdots (s_{m-1}, \gamma_m, a_m, \text{op}_m, R_m, s_m)$ of \mathcal{S} such that $s_m \in F$ is a final state.

Let $\mathcal{V} = (V, \rightarrow, (\curvearrowright^I)_{I \in \mathcal{I}}, \lambda)$ be a simple TCW. Recall that $< = \rightarrow^+$ is the transitive closure of the successor relation. We say that \mathcal{V} is k -round *well timed* with respect to a set of clocks X and stacks $1 \leq s \leq n$ if for each $I \in \mathcal{I}$, the \curvearrowright^I relation for timing constraints can be partitioned as $\curvearrowright^I = \bigsqcup_{1 \leq s \leq n} \curvearrowright^{s \in I} \sqcup \bigsqcup_{x \in X} \curvearrowright^{x \in I}$ where

- (T₁) for each $1 \leq s \leq n$, the relation $\curvearrowright^s = \bigcup_{I \in \mathcal{I}} \curvearrowright^{s \in I}$ corresponds to the matching push-pop events of stack s , hence it is well-nested: for all $i \curvearrowright^s j$ and $i' \curvearrowright^s j'$, if $i < i' < j$ then $i' < j' < j$, see Figure 2.

Moreover, \mathcal{V} consists of at most k rounds, i.e., we have $V = V_1 \sqcup \cdots \sqcup V_k$ with $V_i \times V_j \subseteq <$ for all $1 \leq i < j \leq k$. And each V_ℓ is a round, i.e., $V_\ell = V_\ell^1 \sqcup \cdots \sqcup V_\ell^n$

with $V_\ell^s \times V_\ell^t \subseteq <$ for $1 \leq s < t \leq n$ and push pop events of V_ℓ^s are all on stack s (for all $i \curvearrowright^t j$ with $t \neq s$ we have $i, j \notin V_\ell^s$).

- (T₂) An x -reset block is a maximal consecutive sequence $i_1 < \dots < i_n$ of positions in the domain of the relation $\curvearrowright^x = \bigcup_{I \in \mathcal{I}} \curvearrowright^{x \in I}$. For each $x \in X$, the relation \curvearrowright^x corresponds to the timing constraints for clock x and is well-nested: for all $i \curvearrowright^x j$ and $i' \curvearrowright^x j'$, if $i < i'$ are in the same x -reset block, then $i < i' < j' < j$. Each guard should be matched with the closest reset block on its left: for all $i \curvearrowright^x j$ and $i' \curvearrowright^x j'$, if $i < i'$ are not in the same x -reset block then $j < i'$.

It is easy to check that the simple TCWs defined by a k -dtMPDA are well-timed, i.e., satisfy the properties above.

We denote by $\text{STCW}(\mathcal{S})$ the set of simple TCWs generated by \mathcal{S} . The language of $\mathcal{L}(\mathcal{S})$ is the set of *realizable* simple TCWs in $\text{STCW}(\mathcal{S})$. Given a bound k on the number of rounds, we denote by $\text{STCW}(\mathcal{S}, k)$ the set of simple TCWs generated by runs of \mathcal{S} using at most k rounds. We let $\mathcal{L}(\mathcal{S}, k)$ be the corresponding language. Given a dtMPDA \mathcal{S} , we show that all simple TCWs in $\text{STCW}(\mathcal{S}, k)$ have bounded split-width. Actually, we will prove a slightly more general result. We first identify some properties satisfied by all simple TCWs generated by a dtMPDA, then we show that all simple TCWs satisfying these properties have bounded split-width. Considering $Y = X \cup \{\zeta\}$,

Lemma 6.1. *A k -round well-timed simple TCW has split-width at most $(4nk + 4)(|Y| + 1)$, where n is the number of stacks.*

Proof. Again, the idea is to play the split-game between *Adam* and *Eve*. *Eve* should have a strategy to disconnect the word without introducing more than $(4nk + 4)(|Y| + 1)$ blocks. The strategy of *Eve* is as follows: Given the k -round word w , *Eve* first breaks this into two words. The first word only has stack 1 edges, and the second word has stack edges corresponding to stacks 2, \dots , n . The first word can now be dealt with as we did in the case of TPDA. *Eve* then breaks the second word into two words, the first of which has only stack 2 edges, while the second word has edges of stacks 3, 4 \dots , n , and so on. Finally, we obtain n split-STCW's, each having edges corresponding to only one stack. Once this is achieved, these words can be processed as was done in the case of TPDA. The only thing to calculate is the number of cuts required in isolating each word, which we do in the rest of the proof below.

Obtaining a word containing only stack i edges. Since we are dealing with k -round TCWs, we know that the stack operations follow a nice order : stacks 1, \dots , n are operated in order k times. More precisely, by (T₁), the set V of points of the simple TCW \mathcal{V} can be partitioned into $V = V_1 \cup V'_1 \cup \dots \cup V_k \cup V'_k$ such that $V_1 < V'_1 < \dots < V_k < V'_k$ and $V_1 \cup \dots \cup V_k$ contains all and only stack operations from stack 1. *Eve*'s strategy is to separate all these sets, i.e., cut just after V_1, V'_1, \dots, V_k . This results in $2k - 1$ cuts. This will not disconnect the word if there are edges with clock timing constraints across blocks, i.e., from V_i or V'_i to some other block on the right.

Consider the case when we have timing constraints on some clock x which is reset in some block V_i and checked in some block V'_j with $i \leq j$. All these timing constraints come from the last reset block R_x of clock x which lies in V_i . With two cuts, we detach the consecutive sequence of resets in R_x which are checked in V'_j . So with at most $2|Y|$ cuts, we detach all the resets that come from blocks V_i with $i \leq j$ and are checked in block V'_j .

Similarly, with at most $2|Y|$ cuts, we detach all the resets that come from blocks V'_i with $i < j$ and are checked in block V_j .

In total, including the $2k - 1$ cuts used to separate the blocks $V_1, V'_1, \dots, V_k, V'_k$ we have used $(2|Y| + 1)(2k - 1) \leq 4k(|Y| + 1)$ cuts.

The split-TCW is now disconnected, and we obtain 2 words, one containing only stack 1 edges, and the other containing only stack 2, \dots , n edges. Each word has at most $k|Y|$ holes after the disconnect: at most $|Y|$ holes in each V_i of the first word, and at most $|Y|$ holes in each V'_i of the second word. If $X = \{x_1, \dots, x_m\}$, then the first and second words look like

$$\begin{array}{ccccccc} \overline{V}_1 & H_{x_{g_1}} \dots H_{x_{g_m}} & \overline{V}_2 & H_{x_{d_1}} \dots H_{x_{d_m}} & \dots & \overline{V}_k \\ G_{x_{i_1}} \dots G_{x_{i_m}} & \overline{V}'_1 & G_{x_{h_1}} \dots G_{x_{h_m}} & \overline{V}'_2 & \dots & G_{x_{j_1}} \dots G_{x_{j_m}} & \overline{V}'_k \end{array}$$

Here, $G_{x_{i_j}}$ consists of the bunches of edges coming from the last x_{i_j} -reset block in V_1 , while $H_{x_{g_j}}$ consists of the bunches of edges coming from the last x_{g_j} -reset block in V'_1 and so on. \overline{V}_1 is the word containing the holes which earlier contained $G_{x_{i_1}} \dots G_{x_{i_m}}$ in V_1 and so on. Likewise, \overline{V}'_1 is the word containing the holes which earlier contained $H_{x_{g_1}} \dots H_{x_{g_m}}$ in V'_1 and so on.

Starting with the second word having at most $k|Y|$ holes, we repeat the process to separate out stack 2 edges. This will result again in at most $4k(|Y| + 1)$ edges to be cut, and will result in two new words, each having at most $2k|Y|$ holes. Continuing this, after the last separation, when we isolate stack n edges, we will obtain two words each having at most $nk|Y|$ holes.

Thus, by cutting at most $4nk(|Y| + 1)$ edges, we can separate out the starting word into n words, each having only one kind of stack edge. Applying the TPDA game on each of these words, we obtain a split bound of $(4nk + 4)(|Y| + 1)$. \square

Having established a bound on the split-width for dtMPDA restricted to k rounds, we now discuss the construction of tree automata $\mathcal{B}_{\text{valid}}^{K,M}$ and $\mathcal{A}_S^{K,M}$ when the underlying system is a dtMPDA. As a first step, we keep track of the current round (context) number in the finite control. This makes sure that the tree automaton only accepts runs using at most k -rounds. The validity checks (and hence, $\mathcal{B}_{\text{valid}}^{K,M}$) is as before. The only change pertains to the automaton that checks correctness of the underlying run, namely, $\mathcal{A}_S^{K,M}$, as we need to handle n stacks (and k rounds) instead of the single stack.

Tree Automata Construction for MultiStack and Complexity. Given a dtMPDA $\mathcal{S} = (S, \Sigma, \Gamma, X, s_0, F, \Delta)$, we first construct a dtMPDA \mathcal{S}' that only accepts runs using at most k -rounds. The idea behind constructing \mathcal{S}' is to easily keep track of the k -rounds by remembering in the finite control of \mathcal{S}' , the current round number and context number. The initial states of \mathcal{S}' is $(s_0, 1, 1)$. Here 1,1 signifies that we are in round 1, and context 1 in which operations on stack 1 are allowed without changing context. The states of \mathcal{S}' are $\{(s, i, j) \mid 1 \leq i \leq k, 1 \leq j \leq n, s \in L\}$.

Assuming that $(s, \gamma, a, \text{op}, R, s') \in \Delta$ is a transition of \mathcal{S} , then the transitions Δ' of \mathcal{S}' are as follows:

- (1) $((s, i, j), \gamma, a, \text{op}, R, (s', i, j)) \in \Delta'$ if op is one of nop , \downarrow_c^j or $\uparrow_{c \in I}^j$,
- (2) $((s, i, j), \gamma, a, \text{op}, R, (s', i, h)) \in \Delta'$ if $j < h$ and op is one of \downarrow_c^h or $\uparrow_{c \in I}^h$,
- (3) $((s, i, j), \gamma, a, \text{op}, R, (s', i + 1, h)) \in \Delta'$ if $h < j$ and op is one of \downarrow_c^h or $\uparrow_{c \in I}^h$.

The final states of \mathcal{S}' are of the form $\{(s, i, j) \mid s \in F\}$. It can be shown easily that accepting runs of \mathcal{S}' correspond to accepting k -round bounded runs of \mathcal{S} .

Now, given the dtMPDA \mathcal{S}' , the tree automaton $\mathcal{B}_{\text{valid}}^{K,M}$ that checks for validity and realizability are exactly as in Theorem 4.2.

The only change pertains to the automaton that checks correctness of the underlying run. The tree automaton for the underlying system, $\mathcal{A}_S^{K,M}$ stores the set E^s of pairs (i, j) of left endpoints of *distinct* blocks such that there is at least one push-pop edge pertaining to stack s from $\text{Block}(i)$ to $\text{Block}(j)$. The other change is in the transitions of $\mathcal{A}_S^{K,M}$: the transitions $\text{Add}_{i,j}^I((i, a^1) \oplus (j, a^2))$ in Table 4 will now differ, since this checks the correctness of well-nesting of stack edges when we have only one stack. This transition will be replaced with $\text{Add}_{i,j}^I, s((i, a^1) \oplus (j, a^2))$, which also talks about the particular stack s . Thus, instead of one set E as in the case of TPDA, we require n distinct sets to ensure well-nesting for each stack edge. Secondly, we need to ensure that the k -round property is satisfied. Rather than doing this at the tree automaton level, we have done it at the dtMPDA level itself, by checking this in \mathcal{S}' . This blows up the number of locations by nk , the number of stacks and rounds. Thus, the number of states of the tree automaton $\mathcal{A}_S^{K,M}$ that checks correctness when the underlying system is a k -round dtMPDA is hence $\leq (nk|\mathcal{S}|)^{\mathcal{O}(K)} \cdot 2^{\mathcal{O}(nK^2(|X|+1))}$, where $K = (4nk + 4)(|X| + 2)$.

Proposition 6.2. *Let \mathcal{S} be a k -round multistack timed automaton of size $|\mathcal{S}|$ (constants encoded in unary) with n stacks and set of clocks X . Then, we can build a tree automaton $\mathcal{A}_S^{K,M}$ of size $(nk|\mathcal{S}|)^{\mathcal{O}(K)} \cdot 2^{\mathcal{O}(nK^2(|X|+1))}$ such that $\mathcal{L}(\mathcal{A}_S^{K,M}) = \{\tau \in \mathcal{L}(\mathcal{A}_{\text{valid}}^{K,M}) \mid \llbracket \tau \rrbracket \in \text{STCW}(\mathcal{S})\}$.*

Using arguments similar to Theorem 5.3, we obtain

Theorem 6.3. *Checking emptiness for k -round dtMPDA is decidable in EXPTIME.*

7. CONCLUSION

The main contribution of this paper is a technique for analyzing timed systems via tree automata. This is a new approach, which is quite different from all existing approaches, most of which go via the region/zone based techniques for timed automata. The hardest part of our approach, i.e., checking realizability is oblivious of the underlying timed system and hence we believe that this technique can be applied uniformly to a variety of timed systems. While, we have made a few simplifying assumptions to best illustrate this method, our technique can easily be adapted to remove many of these. For instance, diagonal constraints of the form $x - y \in I$ can be handled easily by adding matching edges. For a constraint $x - y \in [1, 5]$, we add an edge between the last reset of x and last reset of y with $[1, 5]$ interval. Thus, we can check the diagonal constraint at the time of last reset.

As futurework, we would also like to extend our results to other restrictions for dtMPDA such as bounded scope and phase. This would only require us to prove a bound on the split-width and modify the system automaton appropriately to handle the abstract behaviors generated by such systems. Our techniques could also be applied to the more general model [13] of recursive hybrid automata. Another interesting future work is to use our technique to go beyond reachability and show results on model checking for timed systems. While model-checking against untimed specifications is easy to obtain with our approach, the challenge is to extend it to timed specifications. We also see a strong potential to investigate the

emptiness problem for classes of alternating timed automata and hybrid automata. At the very least, we would obtain new width-based under-approximations which may lead to new decidability results.

REFERENCES

- [1] P. Abdulla, M. F. Atig, and J. Stenman. Dense-timed pushdown automata. In *LICS Proceedings*, pages 35–44, 2012.
- [2] C. Aiswarya and P. Gastin. Reasoning about distributed systems: WYSIWYG (invited talk). In *FSTTCS Proceedings*, pages 11–30, 2014.
- [3] S. Akshay, P. Gastin, and S. Krishna. Analyzing timed systems using tree automata. In *CONCUR Proceedings*, 2016.
- [4] R. Alur and D. Dill. A theory of timed automata. In *TCS*, 126(2):183–235, 1994.
- [5] M. F. Atig. Model-checking of ordered multi-pushdown automata. *LMCS*, 8(3), 2012.
- [6] G. Behrmann, A. David, and K. G. Larsen. A tutorial on Uppaal. In *International School on Formal Methods for the Design of Computer, Communication and Software Systems*, pages 200–236, 2004.
- [7] L. Clemente and S. Lasota. Timed pushdown automata revisited. In *LICS Proceedings*, pages 738–749, 2015.
- [8] B. Courcelle. Special tree-width and the verification of monadic second-order graph properties. In *FSTTCS Proceedings*, pages 13–29, 2010.
- [9] B. Courcelle and J. Engelfriet. *Graph Structure and Monadic Second-Order Logic - A Language-Theoretic Approach*, volume 138 of *Encyclopedia of mathematics and its applications*. CUP, 2012.
- [10] A. Cyriac. *Verification of Communicating Recursive Programs via Split-width*. Thèse de doctorat, LSV, ENS Cachan, January 2014.
- [11] A. Cyriac, P. Gastin, and K. Narayan Kumar. MSO decidability of multi-pushdown systems via split-width. In *CONCUR Proceedings*, pages 547–561, 2012.
- [12] W. Czerwinski, P. Hofman, and S. Lasota. Reachability problem for weak multi-pushdown automata. In *CONCUR Proceedings*, pages 53–68, 2012.
- [13] S. N. Krishna, L. Manasa, and A. Trivedi. What’s decidable about recursive hybrid automata? In *HSCC Proceedings*, pages 31–40, 2015.
- [14] S. La Torre, P. Madhusudan, and G. Parlato. The language theory of bounded context-switching. In *LATIN Proceedings*, pages 96–107, 2010.
- [15] S. La Torre, M. Napoli, and G. Parlato. Scope-bounded pushdown languages. In *DLT Proceedings*, pages 116–128, 2014.
- [16] S. La Torre, M. Napoli, and G. Parlato. A unifying approach for multistack pushdown automata. In *MFCS Proceedings*, pages 377–389, 2014.
- [17] P. Madhusudan and G. Parlato. The tree width of auxiliary storage. In *POPL Proceedings*, pages 283–294, 2011.